Modelling the Volatility of Financial Assets using the Normal Inverse Gaussian distribution: With an application to Option Pricing*

Lars Stentoft
Department of Finance
HEC Montréal
Email: lars.stentoft@hec.ca

April 18, 2006

Abstract

In the original GARCH model of Bollerslev (1986) conditional normality is assumed, and although the model was immediately applied to model financial returns the assumption has been found to be violated empirically for many time series. Thus, since then alternative distributions like the Student's t-distribution in Bollerslev (1987) and the Generalized Error Distribution in Nelson (1991) have been introduced and analyzed along with a host of other more or less exotic distributions (see e.g. Lambert & Laurent (2001)). In the present paper we will analyze the Normal Inverse Gaussian distribution, which has been introduced to financial econometrics recently. In relation to the GARCH literature it was used in Forsberg & Bollerslev (2002) to model the ECU/USD exchange rate. In this paper the focus is on modelling individual stock returns, and using Realized Volatility as a measure of daily variance we provide evidence that the Normal Inverse Gaussian framework can provide a good description of the returns of some major US stocks. We also provide estimation results for a NIG GARCH model for these assets. The second part of this paper shows how the NIG GARCH model can be used in an option valuation context. We perform a Monte Carlo study which illustrates that incorporating excess kurtosis and skewness in the option pricing model could potentially explain some of the systematic pricing errors of Gaussian GARCH models found in the literature (see e.g. Stentoft (2005)). However, the results on the empirical performance of the option pricing model are somewhat disappointing and further research is called for in this area.

JEL Classification: C22, C53, G13

Keywords: GARCH models, Normal Inverse Gaussian distribution, American Options, Least Squares Monte-Carlo.

*The author thanks participants at the 3rd Nordic Econometric Meeting as well as seminar participants at HEC Montréal and Brock University. Financial support from SSF (Grant 95015328) is gratefully appreciated.

Contents

1 Introduction

Since the ARCH model was put forth by Engle (1982) and later generalized in Bollerslev (1986) the framework has been used extensively in financial econometrics (see e.g. the surveys by Bollerslev, Chou & Kroner (1992) and Poon & Granger (2003)). However, despite of this the assumptions underlying the original GARCH model of Bollerslev (1986) have been criticized. In particular the assumption of conditional normality has been found not to hold in empirical work. As a result of this, alternative distributions have been suggested and analyzed. The best known examples are probably the Student's t-distribution in Bollerslev (1987) and the Generalized Error Distribution in Nelson (1991), both of which can potentially account for the excess kurtosis often found in the standardized residuals. A general finding in the studies using these alternative distributions is that financial return distributions have fatter tails than the Gaussian distribution. In addition to the leptokurtic distributions various skewed distributions, like e.g. the Skewed Student t-distribution, have been used (see e.g. Lambert & Laurent (2001) which also lists a number of different contributions in this area).

More recently the Normal Inverse Gaussian distribution has been introduced and used in conjunction with GARCH models (see Andersson (2001), Barndorff-Nielsen (1997) and Jensen & Lunde (2001)). Whereas the previously mentioned extensions to the GARCH models have been of an "ad hoc" nature and the choice of distribution has been driven by observed characteristics for the conditional distribution of the residuals, an appealing property of the NIG distribution is that its use can be motivated through a Mixing Distribution Hypothesis (see Clark (1973)). Thus, if returns are assumed to be normally distributed conditional on the variance, and if the conditional variance of the returns follows an Inverse Gaussian distribution, then the unconditional distribution of the returns is a Normal Inverse Gaussian (NIG) distribution.

Thus, an interesting feature with the NIG modelling framework is that it essentially implies a certain behavior for the volatility, although this is in principle unobserved. In the paper by Forsberg & Bollerslev (2002) these implications are analyzed for the ECU/USD exchange rate using the Realized Volatility of Andersen, Bollerslev, Diebold & Labys (2001) as a measure of the daily variance. The findings in that paper are encouraging for this empirically driven approach for the particular exchange rate. In this paper we build on that work extending the ideas to stock returns. As opposed to the exchange rate data our return data does not only have fatter tails than the Gaussian. In addition the distribution is also skewed. Thus, in our formulation of a NIG GARCH model skewness is allowed for, and we provide estimation results for the models which indicate that the extensions to the Gaussian model are important when modelling stock returns.

The present paper also applies the NIG GARCH model to the area of option pricing. Previously, the GARCH models with Gaussian errors have been found to perform well in an option pricing context and compared to the Constant Volatility (CV) models allowing for time-varying volatility has been shown to be important empirically (see Stentoft (2005)). However, the analysis also pointed towards some systematic pricing errors of the Gaussian GARCH model. Thus, in this paper we analyze the possible gains from allowing the returns to be leptokurtic and skewed. Since no theoretical pricing formula exists this is done through a Monte Carlo study. The study shows that by introducing distributions with excess kurtosis and skewness the systematic pricing errors of the Gaussian GARCH model could potentially be explained using the option pricing model of Duan (1999). Empirically, however, the results are somewhat disappointing. In particular, when the model is used on options on three major US stocks it is difficult to find any gains in terms of empirical option pricing errors from having a more flexible distribution for the innovations.

The rest of the paper is organized as follows: As a motivation, in Section 2 the Realized Volatility is used as a measure of daily volatility and we show that the "empirical regularities" found for the ECU/USD exchange rate series in Forsberg & Bollerslev (2002) carry over to the returns for three major US stocks. In Section 3 we present the discrete time framework we will be considering and based on the observations in Section 2 we present a version of the GARCH model with NIG innovations and a parameterization allowing for skewness and excess kurtosis. The section also reports estimation results. Using the option pricing model of Duan (1999), Section 4 describes how the model may be implemented using simulation techniques and in Section 5 a Monte Carlo study is performed to examine the potential gains compared to the Gaussian GARCH benchmark in terms of option pricing. Section 6 describes how we implement the model on actual data and presents the results of our empirical application. Section 7 offers concluding remarks and directions for future research

2 Empirical evidence on Realized Volatility

In our empirical work we will use three major individual stocks from the US. In particular we use General Motors (GM), International Business Machines (IBM), and Merck & Company Inc. (MRK). The reason for choosing these three stocks is in fact that they were highly traded on the options market during the most recent period for which we have data. To be precise, during the period from 1991 through 1995 options on these three stocks were the most traded in terms of actual trades as well as in terms of total volume. Furthermore, return data of good quality for these stocks is readily available from any major data source.

GARCH models were estimated for these particular return series assuming conditional normality in Stentoft (2005), and although that paper found that the GARCH specifications helped explain a number of features found in the data such as volatility clustering and ARCH effects in the standardized residuals, the Gaussian GARCH models failed to explain the excess kurtosis and skewness in the data.

A natural extension to the Gaussian GARCH framework would thus be to allow the error term to follow some alternative distribution which can accommodate these features. The obvious choice would be a leptokurtic skewed distribution and one might think of using the Skewed student t-distribution as it is done in Lambert & Laurent (2001). However, motivated by the work on the ECU/USD exchange rate in Forsberg & Bollerslev (2002) we start out by taking a closer look at a measure of the daily variance called the Realized Volatility, henceforth denoted RV. Our analysis leads us to suggest a GARCH specification with NIG distributed innovations which allows for excess kurtosis as well as skewness.

2.1 Realized Volatility

From the Berkeley Options DataBase (BODB) we can extract not only the relevant option prices but also a series of quotations on the underlying asset as this is reported simultaneously with any quote or trade of the derivatives. For the three stock series under consideration we have observations during the opening hours of the Chicago Board of Options Exchange (CBOE) starting on May 2, 1983 and ending December 29, 1995. With this data we can construct a high frequency data set based on the information that was actually available to the option traders. We follow the procedure outlined in Andersen et al. (2001) of using five minute returns to calculate an estimate of daily volatility. We note that the idea of using sums of squared returns at high frequency to estimate volatility at some lower frequency was used earlier in French, Schwert

& Stambaugh (1987), Hsieh (1991) and Schwert (1989) on daily return data to get estimates of monthly volatilities.

On a typical day, the opening hours at the CBOE were from 9.00 to 16.00 which leaves us with 72 five minute intervals, from the 108'th interval to the 180'th interval measured from midnight. We pick the observation at or immediately before this interval. Since the options we consider are highly traded at the CBOE, the time between the five minute points and the time registered for the observation we use is generally small. Following Andersen et al. (2001) the Realized Volatility is the sum of the squared returns over the intraday intervals. Thus, an estimate of the variance, denoted RV_t^*, can be calculated as

$$RV_t^* = \sum_{i=109}^{180} R_{t,i}^2, \tag{1}$$

where $R_{t,i}$ is the continuously compounded return at day t between interval $i-1$ and i. On a few occasions the Realized Volatility estimated from (1) was extremely low, e.g. with a $\sqrt{RV_t^*} < 0.0005\%$. These observations are likely to indicate days with no trades and thus they were deleted.

It is obvious that the estimate in (1) should not be directly compared to existing measures of interday variance since only a fraction of the day is considered. However, as argued in Hansen & Lunde (2001) there is nothing that suggests that RV_t^* should merely be scaled by the inverse of the fraction of the day considered. The reason is that it is not obvious that an hour with open market should be weighted equally to an hour where the market is closed for volatility calculations. Here we follow the procedure in Hansen & Lunde (2001) of scaling the Realized Volatility such that the sample average equals that of the squared interday returns, R_t^2. As an example, consider the data for General Motors (GM) for which $\frac{\sum R_t^2}{\sum RV_t^*} = 1.226$. Thus, our estimate of daily variance is $RV_t = 1.226 * RV_t^*$.

In the right hand panels of Figure 1 time plots for the $\sqrt{RV_t}$ measure of daily volatility for the series used in this paper are shown. Qualitatively, the pattern that arise is quite similar across the three stocks. In particular, all three plots indicate that volatility is indeed not constant through time and periods of high volatility are followed by low volatility periods and vice versa. This volatility clustering phenomenon corresponds to what is found in the continuously compounded interdaily returns which are shown in the left hand panels in the same figure.

Apart from the clustering of volatility it is clear that the raw returns in the left hand panels are non Gaussian. The left hand plots in Figure 2 show the log density of these series with the Gaussian density superimposed from which it is clear that the unconditional distribution of the raw returns has fatter tails than the Gaussian. This is consistent with previous findings. Panel A of Table 1 shows the unconditional estimation results for the raw return series and, as the tests for departures from normality show, the level of kurtosis in the series significantly exceeds that of the Gaussian distribution. However, a striking feature of the data is that when the returns are standardized by the realized volatility, the fit of the Gaussian distribution is much improved. This can be seen from Panel B of Table 1 showing the sample statistics for the standardized series. Also, the right hand plots of Figure 2 shows the log density plots of the standardized returns, $\frac{R_t}{\sqrt{RV_t}}$, and although departures form normality are still found, these are much less pronounced than for the raw series

2.2 Distributional assumptions on variance

Historically it has been suggested that variance of financial data may be approximated by the lognormal distribution (see e.g. Clark (1973)). However, more recently the Inverse Gaussian distribution, the density of which can be written as

$$f_{IG}\left(z;\sigma^2,a\right)=\frac{\left(\frac{1}{a\sigma^2}\right)^{-\frac{1}{2}}z^{-\frac{3}{2}}}{(2\pi)^{\frac{1}{2}}}\exp\left(a-\frac{1}{2}\left(\frac{a\sigma^2}{z}+\frac{az}{\sigma^2}\right)\right), \tag{2}$$

has been suggested as and alternative to the lognormal assumption, and Barndorff-Nielsen & Shephard (2001) demonstrate that the unconditional distribution of the realized volatility may be well approximated by this distribution. Likewise, in Forsberg & Bollerslev (2002) it is argued that the Inverse Gaussian (IG) distribution is a suitable distribution for the volatility process of the ECU/USD exchange rate.

In Table 2 estimation results are reported on two models for RV. Panel A of the table presents the results for the simplest possible model with a constant, and thus constant volatility, whereas Panel B allows for a ARMA(2,1) specification which means that volatility becomes time varying. In Figure 3 the residual densities are plotted together with those implied by the model.

From Panel A of the table it is clear that the simplest specification corresponding to constant volatility is an insufficient description of the RV process. Although the parameters are estimated with appropriate size and quite precisely in this model, the various test statistics for misspecification are significant for all the RV series. This is particularly so for the $Q(20)$ statistics, which indicates the presence of serial correlation in the standardized residuals. The left hand plots in the figure also points to some shortcomings in this very simple model as the residual plot is quite far from that implied by the estimates.

The results in Panel B from an ARMA(2,1) specification of RV indicates that with a more general specification better results are obtained for the three series. In addition to the ARMA(2,1) results reported here, ARMA(1,1) and ARMA(1,2) models were also estimated. However, the preferred model is in all cases a ARMA(2,1) model.[1] Panel B of the table shows that all the parameters in the ARMA(2,1) model are estimated significantly different form zero, and when the test statistics for misspecification are considered we see that none of these are significant. In the right hand panels of Figure 3 the residual plots are shown for this model, and it is clear from these that the fit by the model is much improved.

Although we do not report the result here the constant model and the ARMA(2,1) models were also estimated with a Log-Normal underlying distribution. In all cases the fit by the Inverse Gaussian model was superior to that of the Log-Normal and this is true both when log-likelihoods are compared and when the Schwartz information criteria is considered. Thus, modelling RV as being Inverse Gaussian distributed provides is an important alternative. However, the results in this section also shows that there is significant correlation in the series. We proposed an ARMA(2,1) model and show that this model provides an important extension to the constant volatility model.

2.3 Distributional assumptions on returns

A particularly nice feature with the IG distribution is that if returns are assumed to be normally distributed conditional on the realized volatility and if the realized volatility follows an Inverse Gaussian distribution, then the unconditional distribution of the returns is a Normal Inverse Gaussian (NIG) distribution. As noted

[1] The results are available upon request from the author.

in Forsberg & Bollerslev (2002) if the realized volatility had been assumed to be lognormally distributed the resulting lognormal mixture distribution for the returns is not readily available except in integral form.

Following Jensen & Lunde (2001) the $NIG(a,b,\mu,\delta)$ distribution can be defined in terms of the so-called invariant parameters as

$$f_{NIG}(x;a,b,\mu,\delta)=\frac{a}{\pi\delta}\exp\left(\sqrt{a^2-b^2}+b\frac{(x-\mu)}{\delta}\right)q\left(\frac{x-\mu}{\delta}\right)^{-1}K_1\left(aq\left(\frac{x-\mu}{\delta}\right)\right), \tag{3}$$

where $q(z)=\sqrt{1+z^2}$ and $K_1(\cdot)$ is the modified Bessel function of third order and index 1. For the distribution to be well defined we obviously need to ensure that $0\le|b|\le a$ and $0<\delta$. We can interpret a and b as shape parameters with a determining the steepness and b the asymmetry. In particular, for $b=0$ we have a symmetric distribution and with a tending to infinity the Gaussian distribution is obtained in the limit. In (3), μ is a location parameter and δ is a scale parameter.

If we define $\rho=b/a$ the mean and variance of a $NIG(a,b,\mu,\delta)$ distributed variable are given as

$$E(X)=\mu+\frac{\rho\delta}{\sqrt{1-\rho^2}}, \text{ and} \tag{4}$$

$$Var(X)=\frac{\delta^2}{a(1-\rho^2)^{3/4}}. \tag{5}$$

Thus, since the NIG distribution is a location-scale family, meaning that $X\sim NIG(a,b,\mu,\delta)\Leftrightarrow\frac{X-\mu}{\delta}\sim NIG(a,b,0,1)$, a zero mean and unit variance Normal Inverse Gaussian variable can be obtained by setting

$$\mu=\frac{-\rho\delta}{\sqrt{1-\rho^2}}, \text{ and} \tag{6}$$

$$\delta=\sqrt{a(1-\rho^2)^{3/4}}. \tag{7}$$

In the following we will refer to this standardized distribution as the NIG(a,b). In Figure 4 we plot the log densities of the Gaussian, the symmetric NIG(a,0) with $a=2$ and $b=0$, and the NIG(a,b) with $a=2$ and $b=0.17$ with these restrictions imposed. The figure illustrates the potential fat tailness and skewness of the NIG distribution.

In Table 3 we report maximum likelihood estimates for the parameters in the NIG distribution along with those from the simple Gaussian model for the sample of returns to be used later. The table clearly shows that the NIG distribution provides a much better fit than does the Gaussian distribution. Furthermore, comparing the two specifications of the NIG distribution, the symmetric NIG(a,0) and the potentially skewed NIG(a,b), we see that skewness is an important feature in our data. We note that when the NIG(a,b) distribution is fitted to a series of returns on the ECU/USD from Datastream corresponding to the one analyzed in Forsberg & Bollerslev (2002), we found the maximum likelihood estimate of b to be -0.0173 with an asymptotic standard error of 0.0332. Thus, for the exchange rate returns we find the skewness parameter to be insignificant which is contrary to what was found for the stock returns.

However, although the analysis above shows the potential for the NIG distribution to fit the leptokurtic unconditional return distributions it is clear that it cannot capture the volatility clustering phenomenon apparent in Figure 1. This is also clear from the test statistics for serial correlation in the squared residuals reported in Table 3. In order to take account of these well documented features the ARCH class of models was combined with the NIG distribution in Barndorff-Nielsen (1997). The class of models has since then been extended to allow for more flexible parameterizations along the lines of the GARCH models in Andersson (2001). In Jensen & Lunde (2001) an even richer volatility structure and asymmetries are allowed for. We now introduce a general framework with leptokurtic and skewed GARCH models which extends all of these previous works.

3 A general framework for leptokurtic and skewed GARCH models

In this paper we consider a discrete time economy with the price of an asset denoted S_t and the dividends of that asset denoted d_t. We will assume that the continuously compounded return process, $R_t = \ln\left(\frac{S_t}{S_{t-1}} + d_t\right)$, can be modelled using the GARCH framework. In the general GARCH model we specify the dynamics for R_t as

$$R_t = m_t\left(\cdot;\theta_m\right) + \sqrt{h_t}\varepsilon_t \tag{8}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h\right) \tag{9}$$

$$\varepsilon_t|\,\mathcal{F}_{t-1} \sim D\left(0, 1; \theta_D\right), \tag{10}$$

where $\mathcal{F}_{t-1}$ is the information set containing all information up to and including time $t-1$.

In equation (8) we use $m_t\left(\cdot;\theta_m\right)$ to denote the conditional mean which is allowed to be governed by the set of parameters θ_m provided that the process is measurable with respect to the information set $\mathcal{F}_{t-1}$. Likewise, in (9) the parameter set θ_h governs the variance process which is allowed to depend on lagged values of innovations to the return process and lagged values of the volatility itself or transformations hereof. Finally, in (10) we use $D\left(0, 1; \theta_D\right)$ to denote some distribution function which is continuous over its support with mean zero and unit variance. The distribution is also allowed to depend on some set of parameters denoted θ_D. For notational convenience we let θ denote the set of all parameters θ_m, θ_h, and θ_D in the following.

To get some intuition, we note that for the well known GARCH(1,1) specification θ_h would consist of the set of parameters $\{\omega, \beta, \alpha\}$ and (9) would correspond to

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}\varepsilon_{t-1}^2. \tag{11}$$

Furthermore, the original Gaussian GARCH model of Bollerslev (1986) is obtained if $D\left(0, 1; \theta_D\right)$ is the standard Gaussian $N\left(0, 1\right)$ distribution, and since the Gaussian distribution is completely characterized by the first two moments no additional parameters are needed in θ_D. Finally, observe that the specification in (8) allows for quite general forms of the conditional mean. In particular, for the GARCH in Mean specification in Duan (1995), which will also be used in this paper, we would have $\theta_m = \{r, \lambda\}$ and

$$m_t\left(\cdot;\theta_m\right) = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t \tag{12}$$

where r is the risk free interest rate. With this specification λ can be interpreted as the risk premium.

3.1 The NIG GARCH model

In the present paper we will impose the restrictions in (6) and (7) on the NIG distribution and use the resulting zero mean and unit variance distribution as the $D\left(0, 1; \theta_D\right)$ distribution in the general GARCH framework in (8) − (10). Thus, $\theta_D = \{a, b\}$. With this specification we have a potentially leptokurtic and skewed GARCH model, which can be interpreted as an extension to the Gaussian GARCH model framework founded in the empirical evidence on Realized Volatility. We note that with $b = 0$ and as a tends to infinity the Gaussian model is obtained in the limit.

Furthermore, we note that if we set $b = 0$ the restrictions in (6) and (7) simplify to $\mu = 0$ and $\delta = \sqrt{a}$. If we further assume that the mean in (8) is zero and the volatility specification in (9) is

$$h_t = \omega + \beta h_{t-1} + \alpha R_{t-1}^2,$$

we essentially obtain the model used for the ECU/USD exchange rate in Forsberg & Bollerslev (2002). Our specification is also similar to that of Jensen & Lunde (2001). In particular, their model implies certain restrictions on the mean, $m_t(\cdot; \theta_m)$, in (8). In our notation we would have the following specifications of (8) and (10)

$$R_t = \mu + \lambda\sqrt{h_t} + \sqrt{h_t}\varepsilon_t, \text{ and} \tag{13}$$

$$\varepsilon_t | \mathcal{F}_{t-1} \sim NIG\left(a, b, \frac{-\rho\delta}{\sqrt{1-\rho^2}}, \sqrt{a\left(1-\rho^2\right)^{3/4}}\right), \tag{14}$$

where $\lambda = \frac{\rho\delta}{\sqrt{1-\rho^2}}$ in (13).

3.1.1 Some important extensions

We now specify a number of extensions to the GARCH model within this framework which can potentially accommodate asymmetric responses to negative and positive return innovations. Thus, these models allow for a leverage effect, which refers to the tendency for changes in stock prices to be negatively correlated with volatility. Some of these particular specifications were used in Stentoft (2005) with assumed Gaussian error distribution.

The first extension to the GARCH model considered is the non-linear asymmetric GARCH model, or NGARCH, of Engle & Ng (1993). The particular specification of the variance process for this model is given by

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}\left(\varepsilon_{t-1} + \gamma\right)^2. \tag{15}$$

Thus, for this specification we would have $\theta_h = \{\omega, \alpha, \beta, \gamma\}$. In the NGARCH model the leverage effect is modelled through the parameter γ, and if $\gamma < 0$ leverage effects are said to be found. It is clear that this model nests the ordinary GARCH specification, which obtains when $\gamma = 0$, and the model thus allows us to compare the contribution of asymmetry directly by comparison to the GARCH specification.

The second extension considered is a version of the exponential GARCH model of Nelson (1991), or EGARCH for short. Instead of (11) the volatility process is now specified in logarithms as

$$\ln(h_t) = \omega + \beta \ln(h_{t-1}) + \alpha\left(|\varepsilon_{t-1}| - E\left[|\varepsilon_{t-1}|\right] + \theta\varepsilon_{t-1}\right). \tag{16}$$

In the EGARCH model the term $|\varepsilon_{t-1}| - E[|\varepsilon_{t-1}|]$ is interpreted as the magnitude effect and $\theta\varepsilon_{t-1}$ is called the sign effect. If $-1 < \theta < 0$ a positive surprise increases volatility less than a negative one, and we say that leverage effects are present. We note that the EGARCH model does not nest the GARCH model. However, it has the nice feature that no restrictions need to be put on the parameters to ensure nonnegativity of the variance.

3.2 Estimation results

Tables 4-6 provide QMLE estimation results from GARCH models with the various volatility specifications in (11), (15), and (16) using return data from 1976 through 1995.[2] The data corresponds to what was used in Stentoft (2005) and consists of time series of continuously compounded returns which have been corrected for dividend payments. The source of the data is the 1997 Stock File from CRSP (see CRSP (1998)). In the estimations the particular mean specification is the one from (12). This is motivated by the results for the option pricing model assuming normality (for more on this see Section 4). As the short rate, r, we take the 1 month LIBOR rate on December 29, 1995, at which date is was 5.4% annualized and continuously compounded. In each table Panel A shows results for the Gaussian case whereas Panel B shows the results for the NIG GARCH model.

Since the results in Panel A of all tables corresponds to those from Stentoft (2005) the we will only discuss the main findings. Thus we note that all the parameters are significant and that the asymmetry parameters in the NGARCH and EGARCH specifications have the expected sign. We also note that while allowing for GARCH volatility specifications is important and does lower the test statistics for ARCH and serial dependency in the standardized returns as well as the squared standardized returns a major problem is that the assumption of normality can be rejected for all series. This is evident from the J-B test statistics which are all significant at any reasonable level. More detailed tests, although not reported here, indicate that it is particularly in terms of excess kurtosis that departures are found although significant skewness is also found to be present in the standardized residuals for a majority of the models. The problem with non normality is also clear from the left hand panels of Figure 5 which plots the log density of the standardized residuals for the Gaussian EGARCH models and clearly indicates fat tails when compared to those from the Gaussian.

In all tables Panel B provide estimation results for the NIG models. These are shown under the relevant headings and indicate the importance of allowing the distribution to depart from the Gaussian. In particular, the tables show large increases in the likelihood values when leptokurtosis and skewness is allowed for when the different GARCH specifications are compared between the Gaussian and the NIG distributions. These increases are shown in parenthesis next to the likelihood value for the NIG models. For all the series the point estimate for a which governs the degree of leptokurtosis indicates a conditional distribution which has quite fat tails. Finally, the b parameter which governs the skewness of the conditional distribution is estimated significantly different from zero for all series. This allows us to conjecture that symmetric although fattailed models like the symmetric NIG model as to simple a model.

When the particular volatility specifications are compared for the NIG conditional distributions the tables show that the leverage effect is important. For all series the γ parameter in the NGARCH specification is estimated significantly different from zero and with the expected sign. For the EGARCH specification the leverage parameter θ is also estimated significantly different from zero and again it has the expected sign. Thus these findings corresponds to those from the Gaussian GARCH models.

In order to be able to compare the different nonnested volatility specifications we report the Schwarz Information Criteria in the last row of each of the panels in the tables. This criteria should be minimized and for all the return series the model achieving the lowest value is the NIG EGARCH model. We note that when only Gaussian models are considered the EGARCH specification is also the preferred one for IBM and MRK although the NGARCH specification minimizes the information criteria for GM as Panel A of Table 4

[2] All computational work was done with Ox (see Doornik (2001)).

shows. Finally, consider the right hand panels of Figure 5 the standardized residuals from the NIG EGARCH are plotted and when compared to the residual plots from the Gaussian EGARCH model we see that the fit is much improved. This conclusion holds for all the returns series and the improvement is particularly pronounced in the tails.

Thus, we end this section by concluding that for the stock return series considered in this paper the leptokurtic and skewed NIG EGARCH model is an important extension to the Gaussian EGARCH model. In particular, the NIG EGARCH model provides a much better fit in the tails of the conditional distribution.

4 Option pricing with leptokurtic and skewed distributions

In Duan (1999) the Local Risk Neutral Valuation Relationship, or LRNVR, of Duan (1995) is generalized to situations where the innovations to the asset return process are non Gaussian. In particular, the Generalized LRNVR, or GLRNVR, applies if this distribution is leptokurtic and skewed. In our setting this corresponds to situations where $D(0,1;\theta_D)$ in (10) departs from the Gaussian distribution.

Under the assumption that the GLRNVR holds, Proposition 4 in Duan (1999) allows us to derive the risk neutralized dynamics of the system in (8) – (10) as

$$R_t = m_t(\cdot;\theta_m) + \sqrt{h_t} F_D^{-1}\left[\Phi\left(Z_t - \lambda_t\right)\right] \tag{17}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h\right) \tag{18}$$

$$\varepsilon_s = F_D^{-1}\left[\Phi\left(Z_s - \lambda_s\right)\right] \text{ if } s \geq t, \tag{19}$$

where Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$, where F_D^{-1} denotes the inverse cumulative distribution function associated with the particular distribution $D(0,1;\theta_D)$, and where Φ denotes the standard Gaussian cumulative distribution function. Thus, when the underlying distribution is the Gaussian, $F_D^{-1}[\Phi(z)] = z$ for any z. However, when the underlying distribution departs from the Gaussian, say by having fatter tails, the transformation yields the innovations under the risk neutral measure with the corresponding properties. Furthermore, in (17) – (19) λ_s is the solution to

$$E^{\mathcal{Q}}\left[\exp\left(m_s(\cdot;\theta_m) + \sqrt{h_s} F_D^{-1}\left[\Phi\left(Z_s - \lambda_s\right)\right]\right)\middle|\, \mathcal{F}_{s-1}\right] = \exp\left(r_s\right), \tag{20}$$

where r_s denotes the one period risk free interest rate at time s. For simplicity, this interest rate is assumed to be deterministic although potentially time varying. Since $m_s(\cdot;\theta_m)$ is measurable with respect to the time $s-1$ information set by assumption, (20) may be rewritten in terms of the "implied" mean function as

$$m_s(\cdot;\theta_m) = r_s - \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_s} F_D^{-1}\left[\Phi\left(Z_s - \lambda_s\right)\right]\right)\middle|\, \mathcal{F}_{s-1}\right]. \tag{21}$$

4.1 The restriction to conditional normality

In the special case of conditional normality $F_D^{-1} = \Phi^{-1}$ and thus $F_D^{-1}[\Phi(\cdot)]$ is the identity function. The restriction in (20) then simplifies to

$$E^{\mathcal{Q}}\left[\exp\left(m_s(\cdot;\theta_m) + \sqrt{h_s}\left(Z_s - \lambda_s\right)\right)\middle|\, \mathcal{F}_{s-1}\right] = \exp\left(r_s\right).$$

Since Z_s is conditional standard normally distributed under the risk neutral probability measure the expectation on the left hand side equals $\exp\left(m_s(\cdot;\theta_m) - \lambda_s\sqrt{h_s} + \frac{1}{2}h_s\right)$. Hence

$$\lambda_s = \frac{m_s(\cdot;\theta_m) - r_s + \frac{1}{2}h_s}{\sqrt{h_s}} \tag{22}$$

is the solution required in (20). With this the system in (17) – (19) simplifies to

$$R_t = r_t - \frac{1}{2}h_t + \sqrt{h_t}Z_t \tag{23}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta\right) \tag{24}$$

$$\varepsilon_s = \left(Z_s - \lambda_s\right) \text{ if } s \geq t, \tag{25}$$

with λ_s given by (22). Note that the mean used in Duan (1995) given by

$$m_t\left(\cdot;\theta_m\right) = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t,$$

where r is the risk free interest rate and λ can be interpreted as the risk premium, leads to the particularly simple expression in (22) of having $\lambda_t = \lambda$.

4.2 GARCH option pricing model using simulation

Following Harrison & Kreps (1979), the time 0 price of any contingent claim should equal its expected payoff under the risk neutral dynamics discounted using the risk free rate of interest. In particular, it follows from the law of iterated expectations, that a European put option with terminal payoff $C_T\left(S_T\right) = \max\left[0, \left(X - S_T\right)\right]$, where X is the strike price and S_T the risk neutralized price of the underlying asset at time T, should have a time 0 price of

$$p_0 = E^{\mathcal{Q}}\left[\exp\left\{-\sum_{t=0}^{T-1} r_t\right\} C_T\left(S_T\right)\middle| \mathcal{F}_0\right], \tag{26}$$

where $E^{\mathcal{Q}}\left[\cdot\middle|\mathcal{F}_0\right]$ means the expectation under the risk neutral measure $\mathcal{Q}$ conditional on the time $t = 0$ information set. In the same manner we can express the time 0 price of an American put option with payoff $C_t\left(S_t\right) = \max\left[0, X - S_t\right]$ if exercised at time t, as

$$P_0 = \sup_{\tau} E^{\mathcal{Q}}\left[\exp\left\{-\sum_{t=0}^{\tau-1} r_t\right\} C_\tau\left(S_\tau\right)\middle| \mathcal{F}_0\right], \tag{27}$$

where the supremum is over all stopping times τ with $0 \leq \tau \leq T$.

An important feature with the system in (17) – (20) is that it specifies the particular risk neutral dynamics to be used above when the underlying model is a GARCH model with a fat tailed and skewed distribution. However, although the pricing system in (17) – (20) is completely self-contained, an actual application of the pricing system to even the simple European option formula in (26) is difficult because of the lack of a closed form expression for the time T value of the underlying asset. Fortunately, it is immediately clear that it is possible to use the system in (17) – (20) to generate a large number of paths with the risk neutralized asset price. From these, an estimated of the option value in (26) can be obtained by taking the average among the paths as

$$\bar{p}_0^M = \frac{1}{M}\sum_{j=1}^{M} \exp\left\{-\sum_{t=0}^{T-1} r_t\right\} \max\left(X - S_T\left(j\right), 0\right), \tag{28}$$

where M is the number of simulated paths and $S_T\left(j\right)$ is the risk neutralized value of the underlying stock at expiration of the option for path number j. Stemming back from Boyle (1977) estimates like (28) have been used to price European options.

Unfortunately, things are not quite as simple when American options are considered. The problem is the need to simultaneously determine the optimal exercise strategy. In particular, simply using

$$\bar{P}_0^M = \frac{1}{M}\sum_{j=1}^{M}\max_{\tau}\left[\exp\left\{-\sum_{t=0}^{\tau-1}r_t\right\}\max\left(X - S_\tau\left(j\right),0\right)\right],$$

to estimate the American put price in (27) would result in a biased estimate, as it is equivalent to assuming that the holder of the option has perfect foresight about future stock prices (see Broadie & Glasserman (1997)). Due to this difficulty, simulation methods were believed to be applicable to the pricing of European style options only until recently (see e.g. Hull (2000, p. 410)). However, the work by Carriere (1996), Longstaff & Schwartz (2001), and Tsitsiklis & Van Roy (2001) has shown otherwise and in these papers algorithms for pricing American style options using simulated values of the underlying asset under the risk neutralized dynamics are proposed.

It is probably fair to say that the most important of these contributions is the Least Squares Monte Carlo (LSM) method of Longstaff & Schwartz (2001) in which expressions like (27) are evaluated using the cross-sectional information available at each step in the simulation. The numerical performance of this method has been examined in Moreno & Navas (2003) and in Stentoft (2004*a*) in the constant volatility Black-Scholes equivalent situation, and recently the mathematical foundation for the use of the method in derivatives research has been provided in Stentoft (2004*b*). Empirically the method is rapidly gaining importance, e.g. for real option valuation (see Gamba (2002)), and recently it was used in Stentoft (2005) to price options in the context of Gaussian GARCH models. In the present paper we will use the algorithm suggested in that paper.

4.3 Implementation of the GARCH option pricing model using simulation

For simplicity, in what follows we will assume that the risk premium parameter, λ_t, and the interest rate, r_t, are constant. With these assumptions and by using (21) to imply the mean, the pricing system is given by

$$R_t = r - \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\Big|\,\mathcal{F}_{t-1}\right] + \sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right], \tag{29}$$

$$h_t = g\left(h_s,\varepsilon_s;-\infty < s \leq t-1,\theta\right), \tag{30}$$

$$\varepsilon_s = F_D^{-1}\left[\Phi\left(Z_s-\lambda\right)\right] \text{ if } s \geq t. \tag{31}$$

From this we see that in addition to generating the standard normally distributed variates Z we also need to be able to calculate the transformation of these variables through $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ and the expectation of the scaled exponential value of this.

Thus, in the following we detail how the method can be applied with particular attention paid to describing how $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ and the logarithm of the expectation $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\big|\,\mathcal{F}_{t-1}\right]$ can be calculated. We examine the relationship for the particular distributions used in this paper and based on this we suggest how these can be efficiently approximated when it comes to implementing the algorithm.

4.3.1 Computing and approximating $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$

Although numerical schemes can be developed to invert the cumulative distribution functions and calculate approximations to $F_D^{-1}\left[x\right]$ following the suggestions in Duan (1999), in the present paper we suggest a different approach. We argue that our procedure is simpler and more flexible. Although the procedure essentially only provides an approximation to $F_D^{-1}[\Phi(Z_t-\lambda)]$ we note that so does the procedure outlined in Duan (1999). The method we propose is based on a random draw from the distribution $D(0,1;\theta_D)$, and consists of the following four steps:

1. First of all a sequence of random normal variates the values of which corresponds to the Z_t's is generated and these are transformed according to $\Phi(Z_t-\lambda)$.
2. Next generate a $1\times m$ vector, y, of independent draws from the distribution under consideration, $D(0,1;\theta_D)$. In most advanced software like Ox this can be done fairly easy with the built in random number generators, RNG's.
3. Calculate which quantile x corresponds to, according to these numbers through the "empirical" distribution function. In Ox the last step can be performed with the command $quantilec(x,y)$. The values from this step corresponds to approximations to $F_D^{-1}[\Phi(Z_t-\lambda)]$ for the Z_t's in step one.
4. Calculate an approximation to $F_D^{-1}[\Phi(Z_t-\lambda)]$, a function of Z, by regressing the values from the step above on a set of transformations of the values in step one.

In Figure 6 the values of $F_D^{-1}[\Phi(Z_t-\lambda)]$ from step three are plotted against $10,000$ random standard normally distributed Z_t from step one with $\lambda=0.05$. Panel A of the figure shows the one to one relation between the Z_t's and $F_D^{-1}[\Phi(Z_t-\lambda)]$ when $D(0,1;\theta_D)$ is the Gaussian distribution. However, the figure also shows that the relationship is smooth and well behaved even when leptokurtic or skewed distributions are used as is indicated by Panel B and C which shows the relationship between the Z_t's and $F_D^{-1}[\Phi(Z_t-\lambda)]$ when the distribution $D(0,1;\theta_D)$ is the symmetric $NIG(2,0)$ and the $NIG(2,0.17)$, respectively.

With Figure 6 in mind we suggest that the function $F_D^{-1}[\Phi(Z_t-\lambda)]$ can be approximated well by a low order polynomial in $Z_t-\lambda$. In our experience good approximations can be obtained with a third order polynomial approximation. This particular choice of approximation is used in the following. Compared to the suggestion in Duan (1999) the use of this approximation will speed up the computational work significantly.

One advantage with the procedure above is that it is not even necessary to know the distribution. Indeed, one could potentially use the actual standardized errors from the estimated model in what we could denote a "bootstrapped" version of the method. Since consistent estimates are obtained when using a QML procedure even though the underlying distribution is actually non Gaussian, this method could in principle allow us to impose the actual empirical distribution on the risk neutral innovations without knowing the particular distributional specification. We note that in Bollerslev & Mikkelsen (1999) the standardized residuals were also used for the simulation. However, in that paper no account was taken of the risk neutralization necessary according to the GLRNVR. The use of the actual standardized residuals has also been suggested by Duan (2002) in a nonparametric framework.

4.3.2 Computing the logarithm of $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}[\Phi(Z_t-\lambda)]\right)\middle|\mathcal{F}_{t-1}\right]$

In Duan (1999) it is suggested that one can calculate the k'th element of the conditional expectation $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t^{(k)}}F_D^{-1}[\Phi(Z_t-\lambda)]\right)\middle|\mathcal{F}_{t-1}\right]$ using "vector" Monte Carlo simulation with N independent standard Gaussian random variables, $\left\{Z_t^i\right\}_{i=1}^N$, for which the inverse is calculated. Thus, the k'th element should be approximated by $\frac{1}{N}\sum_{i=1}^{N}\exp\left(\sqrt{h_t^{(k)}}F_D^{-1}\left[\Phi\left(Z_t^i-\lambda\right)\right]\right)$. However, in our experience this is extremely time consuming. The procedure we suggest exploits the fact that h_t is measurable with respect to the current information set and the logarithm of the expectation, $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda_t\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$, is a relatively simple function $f(h_t)$ which can be easily approximated. To be specific, our procedure consists of the following three steps

1. Select appropriate numbers for h_t. This may either be done by selecting random numbers from an appropriate distribution or by simply choosing a range of potential variances. In particular, if we believe that variances are Inverse Gaussian distributed a random draw from this distribution may be used.

2. Based on the approximation scheme outlined for $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ in steps one through three above, for each h_t calculate the logarithm of the expectation as the logarithm of the mean of the exponential values of the product of the particular h_t and the approximation $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ based on a random draw of Z_t's.

3. Calculate an approximation to $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$, a function now only of h, by regressing the values from the step above on a set of transformations of the values in step one.

In figure 7 we plot approximate values of this function $f(h_t)$ against values of h_t, where the expectation is based on a sample of random standard normally distributed variables Z_t with $\lambda = 0.05$ for which $F_D^{-1}\left[\Phi\left(Z_t-\lambda_t\right)\right]$ have been calculated as outlined above. Panel A of the figure shows that the approximation is very close to the exact expectation for the Gaussian case which we know is given as $f(h_t) = \frac{1}{2}h_t - \lambda\sqrt{h_t}$. The other panels show that the functional relationship remains smooth and well behaved when fat tails and skewness are allowed for.

Based on the shape of the functional relationship we suggest that the function $f(\cdot)$ be approximated with a low order polynomial in h_t. In our experience good approximations can be obtained when only a third order approximation is used. Again we note that compared to the suggestion in Duan (1999) the use of this approximation will speed up the computational work significantly.

5 A Monte Carlo study of the GARCH option pricing model

With the approximation schemes outlined above we are ready to use the algorithm to price American options under GARCH when the underlying distribution departs from the Gaussian. In this section we provide an analysis of the effects on the option prices of these features, and the results are compared to what would be obtained under the assumption of normality using a set of artificial options largely corresponding to those used in Stentoft (2005).

The paper by Stentoft (2005) showed that the early exercise feature is important when a GARCH specification is the correct underlying model and that the pricing errors occurring when a constant volatility model is used incorrectly are consistent with what was found empirically for a number of individual stock options. In particular, smiles in implied Black-Scholes volatilities and underpricing of short term options and out of the money options where shown to obtain.

The purpose of the present Monte Carlo study is to examine the effect of allowing for conditional excess kurtosis and skewness, thus generalizing these results further. We note that some results on the matter are available in Duan (1999) where the Generalized Error Distribution, or GED for short, is used as an alternative to the Gaussian together with the NGARCH variance specification. We will elaborate on these findings in a number of directions. In particular, in the present study the analysis is performed for American style options, and empirically plausible parameter values are used. Last but certainly not least, we introduce skewness through the NIG GARCH specification.

5.1 The Monte Carlo setup

To illustrate the pricing effects we use a set of artificial options with strike prices ranging from deep out of the money to deep in the money, and with ultra short maturity (1 week), short maturity (1 month), middle maturity (3 months) and long maturity (6 months). We take a year to be 252 trading days, such that the options expire in 7, 21, 63, and 126 trading days respectively. To a large extent this collection of options covers what is actually observed for traded options on individual stocks. Based on the analysis for the Gaussian situation in Stentoft (2005) in the following we will report results only for the GARCH and the NGARCH variance specifications.

When comparing different volatility specifications within the GARCH framework it seems reasonable that we consider models with the same characteristics at least in terms of the level of persistence and the level of implied volatility. We do this by first choosing the persistence parameters, corresponding to $\beta + \alpha$ in the GARCH model and $\beta + \alpha\left(1+\gamma^2\right)$ in the NGARCH model, and then adjusting ω in order to yield the same unconditional variance, $E[h]$, across the models. Thus, in the GARCH model we set $\omega_{GARCH} = (1-\beta-\alpha) * E[h]$ and in the NGARCH model we set $\omega_{NGARCH} = \left(1-\beta-\alpha\left(1+\gamma^2\right)\right) * E[h]$.

A second request in our Monte Carlo study is that the parameter values for λ, α, β, γ, and $E[h]$ be empirically plausible and again we choose values close to the actual averages of the estimated parameter values from the individual return series from 1976 through 1995 reported in Chapter 3.2. After rounding we decide to use $\beta = 0.92$ and $\alpha = 0.06$ in the GARCH specification, and $\beta = 0.92$, $\alpha = 0.048$, and $\gamma = -0.5$ in the NGARCH model. These values yield a persistence of 0.98 in both models, so in order to get an annualized unconditional volatility of 25% we set $\omega = 4.96 \times 10^{-6}$. Finally, we set $\lambda = 0.05$ in all the models. In order to start up the simulations we need either values for h_0 and ε_0 for the models or we can specify h_1. We choose the latter, as does Duan (1995) and Stentoft (2005), and set h_1 equal to the unconditional expectation, that is $E[h]$. Finally, in line with what is found empirically we set the interest rate equal to 6% on an annual basis. Furthermore, corresponding to what is found for our subset of individual stocks we choose to include dividend payments and we choose a rate of 3% annually.

5.2 Pricing American options under GARCH with leptokurtic and skewed distribution

Table 7 presents the results for the GARCH variance specification with

$$\begin{aligned} S_t &= S_{t-1} \exp\left\{(0.06-0.03)/252 - \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t} F_D^{-1}\left[\Phi\left(Z_t - 0.05\right)\right]\right)\right] + \sqrt{h_t} F_D^{-1}\left[\Phi\left(Z_t - 0.05\right)\right]\right\} \\ h_t &= 4.96 \times 10^{-6} + 0.92 * h_{t-1} + 0.06 * h_{t-1} F_D^{-1}\left[\Phi\left(Z_{t-1} - 0.05\right)\right]^2, \end{aligned}$$

where $Z_t \sim N(0,1)$ and where the quantities $F_D^{-1}\left[\Phi\left(Z_t - 0.05\right)\right]$ and $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t} F_D^{-1}\left[\Phi\left(Z_t - \lambda\right)\right]\right)\right]$ are approximated according to the algorithms described above. Table 8 presents the results for the NGARCH specification using the same mean equation but with a specification for the variance process corresponding to

$$h_t = 4.96 \times 10^{-6} + 0.92 * h_{t-1} + 0.048 * h_{t-1} \left(F_D^{-1} \left[\Phi \left(Z_{t-1} - 0.05 \right) \right] - 0.5 \right)^2 .$$

The reported prices are averages of 100 calculated estimates using different seeds in the random number generator. In the table the standard errors of these 100 estimates are reported next to the corresponding price estimate. For the price estimates products and cross-products between the stock level and the volatility level of total order less than or equal to two are used in the cross-sectional regressions and in all the cases exercise is considered once every day. In column three and four of both panels of both tables the estimates under conditional Gaussian innovations are reported. These values correspond to what was reported in Stentoft (2005). The rest of the table now generalizes this result by comparing the Gaussian results to the results obtained with alternative distributional assumption.

Column five and six of Panel A in Table 7 reports the American put price estimates and the standard errors obtained for a symmetric NIG distribution with $a = 2$ and $b = 0$, whereas column seven reports the difference between the Gaussian GARCH estimate and the symmetric NIG GARCH estimate relative to the latter. Thus, this column, which is headed "Rel Diff", can be interpreted as the mispricing that would occur if a Gaussian GARCH model was used when the true model was in fact a NIG GARCH model with $NIG(2, 0)$ distributed errors. The table shows that compared to the symmetric NIG GARCH model the Gaussian GARCH model severely underprices the short term out of the money options. On the other hand, the at the money options are slightly overpriced by the Gaussian GARCH model. Turning the attention to the corresponding columns in Panel B of the table we see that equivalent results are found for the estimated American call prices.

Column eight of Panels A and B in Table 7 reports the American option price estimates when skewness in the distribution is allowed for in the NIG distribution by setting $b = 0.2$ and column ten compares these with the Gaussian GARCH model for the puts and calls respectively. The latter column shows that, although the underpricing of the out of the money options is smaller than for the symmetric $NIG(2, 0)$ for the put options in Panel A it is much larger for the call options in Panel B. Furthermore, for the call options considered the mispricing persist and is large in economic terms even for the long term out of the money options.

For the NGARCH model, Table 8 shows that the overall results are in line with those for the GARCH specification. However, if the volatility specification has asymmetries the amount of mispricing by the Gaussian model is reduced somewhat for the put options relative to the various alternative distributions, whereas it is increased significantly for the call options. This holds particularly for the symmetric NIG GARCH model.

In Figures 8 and 9 the implied volatility backed out from the constant volatility Binomial Model is plotted against the corresponding difference between the strike price and the level of the stock. Although these plots convey the same general message as the tables, the various plots indicated that large differences in terms of pricing can be expected to be found in particular for the short term options. This is illustrated well by Figure 9 where the top two panels indicates that large price differences can occur for short term out of the money call options when the conditional distribution is allowed to depart from the Gaussian.

6 Implementing the GARCH option pricing model

Our sample of options consists of options on the three stocks GM, IBM, and MRK. We use weekly data from 1991 through 1995 and sample options on each Wednesday. If no options are traded on the Wednesday we pick the Thursday immediately after. At any particular day we sample an end of day option price for all contracts, that is combinations of strike price and maturity, for which the traded volume during the day was at least five. In our empirical application we assume that the options can be exercised only once a day. We use a total of $M = 20,000$ and in estimating the conditional expectations in the LSM algorithm, apart from a constant term, we use powers of and cross products between the asset price and the level of the variance of total order less than or equal to two. We note that although term structure effects could be included in the present paper we assume a constant interest rate. Furthermore, we treat cash dividend payments as known both in size and timing and let the payment spill over on the stock level fully. For more on this see also Stentoft (2005).

Under the assumption of conditional lognormality the method can be immediately applied empirically as this is done in Stentoft (2005). In particularly, with the specific choice of mean specification of Duan (1995) the restriction on λ_t becomes particularly simple as we have already seen. This greatly simplifies both the estimation and the simulation procedure. However, when the conditional distribution is non Gaussian it is generally not possible to find as simple an expression for λ_t. One possible alternative is to "imply" the mean $m_t(.;\theta_m)$ in (17) by replacing it with $r_t - \ln E^{Q}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda_t\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ as it was suggested above. However, one should note that for consistency the same then has to be done in (8) which in terms alters the specification of the model used for estimation.

It is immediately clear that this potentially complicates the estimation procedure since this expression would have to be evaluated numerically in search of maximum likelihood estimates. In fact, in the numerical part of Duan (1999) it is argued that this procedure becomes extremely demanding in computing time and instead the above problem is dealt with by assuming constant interest rates, a constant value for λ, and restricting the mean to be a constant. However, it is noted that the parameter λ is in fact left unidentified by such a procedure.

Here we argue that the procedure of implying the solution to (20) is feasible and that the estimation procedure does not necessarily become computationally excessively demanding since this expectation can be calculated easily following the procedures above. However, the effect on the parameter estimates of this is limited and can be safely neglected. To be specific, consider a specification with constant $\lambda_{t+1} = \lambda$, and assume for the time being that the interest is zero. Then the mean $m_{t+1}(.;\theta_m)$ in the estimation step should be replaced by $-\ln E^{Q}\left[\exp\left(\sqrt{h_{t+1}}F_D^{-1}\left[\Phi\left(Z_{t+1}-\lambda\right)\right]\right)\middle|\mathcal{F}_t\right]$, and we know from the discussion above that this expectation can be approximated numerically in a simple way. In Table 9, Panel A reports the results of implying the parameter for GARCH(1,1) specification with Normal Inverse Gaussian errors whereas Panel B reports the results from the ordinary estimation procedure. A striking feature is that virtually no difference exists between the parameter estimates. Thus, since the "regular" estimation procedure is more time efficient we use this in what follows and turn to the empirical results.

6.1 Review of previous findings

In Stentoft (2005) the Gaussian GARCH model was used to price American style options on the same three US stocks we have considered in the present paper. The results from that paper were that although introducing time-varying volatility into option pricing decreased the pricing errors, the remaining errors could still in part be explained by particularly the moneyness of the options. To be precise, the regression results in Stentoft (2005) indicate that a smile type relationship between the moneyness of the options and the relative bias of say the GARCH model exists. Although this is not directly comparable to the smile in implied volatilities it still indicates a systematic pricing error which should be incorporated in the theoretical option pricing model, if at all possible.

The relationship found in Stentoft (2005) would indicate that the in the money options are priced relatively lower than the out of the money options. Comparing this to our Monte Carlo results there is reason to believe that allowing for excess kurtosis and potentially also skewness could explain some of the systematic pricing errors found in that paper since this is exactly what is implied by the results in Tables 7 and 8. However, whether or not this extension to the Gaussian GARCH option pricing model suffices to explain the regularities found in Stentoft (2005) is an empirical question to which we turn our attention in the following.

6.2 Overall pricing results for the NIG GARCH models

Table 10 examines the overall pricing results for the NIG GARCH option pricing model using some relevant metrics from the literature. Letting P_k and $\tilde{P}_k$ denote the k'th observed price respectively model price these are the relative mean bias, $RBIAS \equiv K^{-1} \sum_{k=1}^{K} \frac{(\tilde{P}_k - P_k)}{P_k}$, and the relative mean squared error, $RSE \equiv K^{-1} \sum_{k=1}^{K} \frac{(\tilde{P}_k - P_k)^2}{P_k^2}$. The reason for using these relative metrics is that the actual option prices varies significantly between the in the money and out of the money options. Thus, if absolute metrics were used instead much more weight would be put on pricing the in the money options correct. In the table Panel A provides the results when all the options are combined whereas Panel B through D provides the pricing results for the individual companies. Each panel provides pricing results for all options as well as the put options and the call options respectively.

Considering first Panel A of the table it can be observed, as it has been found for the Gaussian GARCH models, that large improvements is found for the NIG GARCH model when compared to the constant volatility (CV) model with NIG errors. Thus, the first conclusion is that allowing the volatility to be time varying is important when pricing American options on the individual stocks considered in this paper. The results hold true when the put options and the call options are considered individually, and the panel shows that the CV model is the one with the largest pricing errors in all cases and using both of the chosen metrics. In addition to this, the panel shows that not only are GARCH features important but so is asymmetries. In particular, either the NGARCH or the EGARCH models are the best performing models in the majority of the cases. In particular this is so for all the cases when the RSE metric is used although the symmetric GARCH model has somewhat smaller errors according to the RBIAS metric for the call options and also slightly smaller pricing errors when considering both puts and calls.

Next consider the individual results in Panels B through D which again shows that using the CV model results in the largest pricing errors irrespective of which underlying stock, which metric, or which type of options are considered. However, the panels also shows that while there is improvements in the pricing performance for the all the stocks the magnitude of the improvements varies. Thus, while Panel B shows that the overall pricing errors of the CV model is about nine times that of either of the GARCH specifications for GM the improvements is only a fraction of this for MRK, with the results of IBM in between. When the models are ranked it turns out that the best performing model is the EGARCH model for the majority of the cases although both the NGARCH and the symmetric GARCH models performs better in some instances. Considering the puts and calls together the results for the RBIAS metric shows that there is generally not much difference in the pricing performance of the models except for IBM where the EGARCH model performs surprisingly poorly. When using the RSE metric the panels shows that when both puts and calls are considered asymmetries appears quite important.

6.3 Pricing performance relative to the Gaussian GARCH model

Table 11 compares the pricing performance of the NIG GARCH model to what is obtained with a Gaussian GARCH model and reports the ratio of the NIG GARCH pricing errors to those from Gaussian GARCH model. Thus, if a cell number is smaller than one it indicates that the NIG model has the smallest pricing error for this metric and variance specification. On the other hand cell numbers larger than one indicates that the Gaussian model actually has the smallest pricing errors. Again Panel A provides the results when all the options are combined whereas Panel B through D provides the pricing results for the individual companies and each panel provides pricing results for all options as well as the put options and the call options respectively.

Considering first Panel A of the table it is observed that the number generally seem to favour the Gaussian GARCH model. In particular, this is so for the RBIAS metric where the ratios are quite high at least when the results for the CV specification are ignored. Although the NGARCH specifications are performing close to equivalently well for the call options for the GARCH and EGARCH specifications the ratios are generally higher than 1.7 - that is the RBIAS pricing errors of the NIG model are 70% larger that what was found for the Gaussian model in virtually all cases!

However, the RBAIS is only one potential metric and when the results for the RSE metric are considered quite different results are obtained. To be specific, when this metric is used the table shows that the Gaussian GARCH models are outperformed by the NIG GARCH models in the majority of cases. In fact this holds for both the GARCH and the NGARCH volatility specification used on put or call or both types of options and for the EGARCH specification when call options alone are considered. Although the improvements in RSE errors are smaller in magnitude than the increase in RBIAS errors in multiple situations the reduction in pricing errors is $15\% - 20\%$ when the NIG models are compared to the corresponding Gaussian models. The table shows that this holds in particular for call options where all NIG GARCH models outperforms the corresponding Gaussian GARCH models.

When Panels B through D are considered it is seen that large differences occur between the individual stock options. This holds particularly for the RBIAS metric and particularly so for options on IBM. For this stock the NIG GARCH and EGARCH models have pricing errors which are more than double what is obtained with the Gaussian GARCH and EGARCH models. The poor performance is seen to stem primarily from the put options. With the exception of the EGARCH specification the results are closer for GM and MRK. However, it may be observed also that the panels shows quite different results when the RSE metric is used and using this metric it is once again found that the NIG models in most cases will have smaller pricing errors than the corresponding Gaussian models. a

6.4 Moneyness and maturity effects

From the Monte Carlo study above we learn that introducing excess kurtosis generates a virtual smile or smirk in the prices from a Gaussian GARCH model. Thus, in this section the pricing results are reported for different categories of moneyness measured as the ratio between the asset price and the discounted strike price, $Mon = S/(K * \exp(-r * T_1))$. As the cut of point we choose 2% and thus in the money options are options with a moneyness in excess of 102% whereas out of the money options have $Mon < 92\%$. We also report results for different categories of maturity, as we partition the sample of options into short term options with less than 42 trading days to maturity, ST, middle term options with between 42 and 126 trading days to maturity, MT, and the long term options with more than 126 trading days to maturity, LT. The results for these two partitions both in absolute terms as well as relative to the Gaussian GARCH models are shown in Table 12.

Panel A shows the overall results in terms of moneyness, and from these panels it is clear that as the options move away from the moneyness the pricing errors increase. This holds both for the RBIAS and the RSE metrics and for all the NIG GARCH models. It remains true though that the CV model is outperformed by the GARCH specifications and in all but one of the partitions an asymmetric model is the best performing one.

In Panel B the overall pricing results of the NIG GARCH models is compared to those from the Gaussian GARCH models and the relative performance is reported, that is the ratio between pricing errors of the NIG model and the corresponding specification with Gaussian errors. From this panel it is clear that there is in fact no worsening of the relative pricing performance as we move away from the moneyness. In fact there is a tendency for the NIG GARCH model to perform better than the Gaussian GARCH models for the OTM option than for the ITM options, and when the RSE metric is used the panel shows that the NIG specifications actually outperforms the Gaussian specifications for OTM options for the GARCH and NGARCH models.

In Panel C the overall results in terms of maturity are shown, and from the panel it is observed that there is no clear relationship between the pricing performance and maturity. In fact while the RBIAS errors increase slightly with maturity the RSE errors decrease. This holds for all volatility specifications. Once the pricing performance is related to that of the Gaussian GARCH models there is in fact a clear tendency for the pricing errors to increase with maturity as Panel D shows, and whereas the RSE errors for the NIG models was 10-20% less for the short maturity options it is 10-50% larger for the long maturity options.

7 Conclusion

In the present paper an empirical foundation for using the NIG GARCH model as a model for stock price returns is provided using the Realized Volatility measure of daily variance. In particular, evidence is provided showing that although the distribution of raw returns departs from the Gaussian, when the returns are standardized with the Realized Volatility the Gaussian density provides a much better approximation. Building on this the NIG distribution is suggested as a relevant distribution for the innovations to the return process. However, in order to take account of the volatility clustering phenomenon found in financial time series, the actual modelling is done in the GARCH framework. We estimate the ensuing model, the NIG GARCH model, on returns from three major US assets and find that the extra flexibility of the NIG GARCH model when compared to the Gaussian GARCH model results in significantly increases in the Likelihood.

Furthermore, the NIG GARCH model is the chosen one when the Schwarz Information Criteria is used.

Having established the empirical foundation for the NIG GARCH model it is used in an option pricing context. Through a Monte Carlo experiment we show that the ability of the NIG GARCH model to accommodate conditional leptokurtosis and skewness can potentially explain some of the systematic pricing errors found in recent empirical work on option pricing with time-varying volatility. In particular, the NIG GARCH model has the ability to generate higher prices than the Gaussian GARCH model for out of the money put and call options particularly of short maturity.

In the paper it is noted that this is typically the options for which Gaussian GARCH models are found empirically to have the largest pricing errors. It is thus of interest to perform an empirical evaluation of the model in order to evaluate its potential in this relation and this is done in the final part of this paper. Unfortunately, from our empirical application using data from 1991 through 1995 on three major US stocks it is difficult to conclude that the added flexibility in the NIG GARCH option pricing model constitutes a valuable extension to the Gaussian GARCH option pricing model.

An obvious idea for a generalization of the method is to use the standardized returns in the simulation as a "Bootstrapped" version. Although our preliminary research has been inconclusive as to the potential of this method the research in this area is ongoing. Furthermore, the analysis in the present paper has used historical return data only for estimation. However, as the results for e.g. IBM showed such data alone might not enable us to predict the level of future volatility. Including historical option data in the analysis could potentially help in this direction and thus potentially provide better price estimates.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2001), 'The Distribution of Exchange Rate Volatility', *Journal of the American Statistical Association* **96**(453), 42–55.

Andersson, J. (2001), 'On the Normal Inverse Gaussian Stochastic Volatility Model', *Journal of Business and Economic Statistics* **19**(1), 44–54.

Barndorff-Nielsen, O. E. (1997), 'Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling', *Scandinavian Journal of Statistics* **24**, 1–13.

Barndorff-Nielsen, O. E. & Shephard, N. (2001), 'Econometric analysis of realized volatility and its use in estimating stochastic volatility models', *Journal of the Royal Statistical Society, Series B* **63**, 167–241.

Bollerslev, T. (1986), 'Generalized Autoregressive Conditional Heteroskedasticity', *Journal of Econometrics* **31**, 307–327.

Bollerslev, T. (1987), 'A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return', *Review of Economics and Statistics* pp. 542–547.

Bollerslev, T., Chou, R. Y. & Kroner, K. F. (1992), 'ARCH Modelling in Finance', *Journal of Econometrics* **52**, 5–59.

Bollerslev, T. & Mikkelsen, H. O. (1999), 'Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics', *Journal of Econometrics* **92**, 75–99.

Boyle, P. P. (1977), 'Options: A Monte Carlo Approach', *Journal of Financial Economics* **4**, 323–338.

Broadie, M. & Glasserman, P. (1997), 'Pricing American-Style Securities Using Simulation', *Journal of Economic Dynamics and Control* **21**, 1323–1352.

Carriere, J. F. (1996), 'Valuation of the Early-Exercise Price for Options Using Simulations and Nonparametric Regression', *Insurance: Mathematics and Economics* **19**, 19–30.

Clark, P. K. (1973), 'A subordinated stochastic process model with finite variance for speculative prices', *Econometrica* **41**(1), 135–155.

CRSP (1998), *CRSP Stock File Guide*, Center for Research in Security Prices, Graduate School of Business, University of Chicago.

Doornik, J. A. (2001), *Ox: An Object-Oriented Matrix Language*, 4th edn, Timberlake Consultants Press, London.

Duan, J.-C. (1995), 'The GARCH Option Pricing Model', *Mathematical Finance* **5**(1), 13–32.

Duan, J.-C. (1999), 'Conditionally Fat-Tailed Distributions and the Volatility Smile in Options', *Mimio, Hong Kong University of Science and Technology* .

Duan, J.-C. (2002), 'Nonparametric Option Pricing by Transformation', *Mimeo, Rotman School of Management, Toronto* .

Engle, R. F. (1982), 'Autoregressive Conditional Heteroscedasticity With Estimates of The Variance of United Kingdom Inflation', *Econometrica* **50**(4), 987–1007.

Engle, R. F. & Ng, V. K. (1993), 'Measuring and Testing the Impact of News on Volatility', *Journal of Finance* **48**(5), 1749–1778.

Forsberg, L. & Bollerslev, T. (2002), 'Bridging the Gap Between the Distribution of Realized (ECU) Volatility and ARCH Modeling (of the EURO): The GARCH Normal Inverse Gaussian Model', *Journal of Applied Econometrics* **17**, 535–548.

French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987), 'Expected Stock Returns and Volatility', *Journal of Financial Economics* **19**, 3–29.

Gamba, A. (2002), 'Real Options Valuation: A Monte Carlo Approach', *Mimeo, University of Verona, Italy* .

Hansen, P. R. & Lunde, A. (2001), 'A Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)?', *Working Paper, Centre for Analytical Finance* **84**.

Harrison, J. M. & Kreps, D. M. (1979), 'Martingales and Arbitrage in Multiperiod Securities Markets', *Journal of Economic Theory* **20**, 381–408.

Hsieh, D. A. (1991), 'Chaos and Nonlinear Dynamics: Applications to Financial Markets', *The Journal of Finance* **46**(5), 1839–1877.

Hull, J. C. (2000), *Options, Futures, and Other Derivatives*, Prentice-Hall, Inc., Upper Saddle River, New Jersey.

Jensen, M. B. & Lunde, A. (2001), 'The NIG-S&ARCH model: a fat-tailed, stochastic, and autoregressive conditional heteroskedastic volatility model', *Econometrics Journal* **4**, 319–342.

Lambert, P. & Laurent, S. (2001), 'Modelling Financial Time Series Using GARCH-Type Models with a Skewed Student Distribution for the Innovations', *Mimeo, Université de Liège* .

Longstaff, F. A. & Schwartz, E. S. (2001), 'Valuing American Options by Simulation: A Simple Least-Squares Approach', *Review of Financial Studies* **14**, 113–147.

Moreno, M. & Navas, J. F. (2003), 'On the Robustness of Least-Squares Monte Carlo (LSM) for Pricing American Options', *Review of Derivatives Research* **6**, 107–128.

Nelson, D. B. (1991), 'Conditional Heteroskedasticity In Asset Returns: A New Approach', *Econometrica* **59**(2), 347–370.

Poon, S.-H. & Granger, C. W. J. (2003), 'Forecasting Volatility in Financial Markets: A Review', *Journal of Economic Literature* **XLI**, 478–539.

Schwert, G. W. (1989), 'Why Does Stock Market Volatility Change Over Time?', *The Journal of Finance* **44**(5), 1115–1153.

Stentoft, L. (2004*a*), 'Assessing the Least Squares Monte-Carlo Approach to American Option Valuation', *Review of Derivatives Research* **7**(3), 129–168.

Stentoft, L. (2004*b*), 'Convergence of the Least Squares Monte Carlo Approach to American Option Valuation', *Management Science* **50**(9), 1193–1203.

Stentoft, L. (2005), 'Pricing American Options When the Underlying Asset follows GARCH Processes', *Journal of Empirical Finance* **12**(4), 576–611.

Tsitsiklis, J. N. & Van Roy, B. (2001), 'Regression Methods for Pricing Complex American-Style Options', *IEEE Transactions on Neural Networks* **12**(4), 694–703.

Tables

Table 1: Estimation results for the unconditional models - the sample statistics

Panel A: Raw Returns

	GM	(#3182)	IBM	(#3177)	MRK	(#3186)
	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
μ	0.0497	(0.0308)	0.0157	(0.0272)	0.1020	(0.0262)
ω	3.0481	(0.2063)	2.4061	(0.2402)	2.1983	(0.0964)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
Skewness	196.67	[0.0000]	1121.12	[0.0000]	23.01	[0.0000]
Ex.Kurt.	20950.74	[0.0000]	116379.54	[0.0000]	2265.39	[0.0000]
$J-B$	21147.41	[0.0000]	117500.67	[0.0000]	2288.41	[0.0000]
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	35.61	[0.0171]	34.14	[0.0252]	33.96	[0.0264]
$Q^2(20)$	421.63	[0.0000]	128.22	[0.0000]	397.15	[0.0000]
$ARCH1-5$	82.84	[0.0000]	21.47	[0.0000]	24.80	[0.0000]

Panel B: Raw Returns

	GM	(#3182)	IBM	(#3177)	MRK	(#3186)
	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
μ	0.0317	(0.0176)	0.0245	(0.0167)	0.0933	(0.0187)
ω	0.9876	(0.0244)	0.8910	(0.0262)	1.1133	(0.0299)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
Skewness	9.22	[0.0024]	0.30	[0.5846]	15.10	[0.0001]
Ex.Kurt.	0.35	[0.5538]	72.19	[0.0000]	12.25	[0.0005]
$J-B$	9.57	[0.0084]	72.49	[0.0000]	27.35	[0.0000]
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	35.13	[0.0194]	19.50	[0.4893]	44.01	[0.0015]
$Q^2(20)$	19.90	[0.4643]	28.13	[0.1064]	96.59	[0.0000]
$ARCH1-5$	1.20	[0.3085]	2.56	[0.0257]	7.09	[0.0000]

Notes: This table reports the sample statististics for the raw returns and the standardized returns. In parentheses the standard errors are reported. In brackets P-values for the relevant test statistics are shown.

Table 2: Estimation results for RV with Inverse Gaussian errors

Panel A: Estimation results with constant model

	GM (#3180)		IBM (#3178)		MRK (#3186)	
Likelihood	-6047.00		-5392.22		-5233.95	
	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
σ^2	3.0098	(0.0709)	2.4177	(0.0956)	2.3539	(0.1366)
a	1.4458	(0.0769)	1.3322	(0.1608)	1.1402	(0.1407)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	3410.860	[0.000]	2603.573	[0.000]	1691.111	[0.000]
$Q^2(20)$	652.631	[0.000]	281.948	[0.000]	647.862	[0.000]
$ARCH1-5$	93.503	[0.000]	55.477	[0.000]	23.863	[0.000]
Schwarz	3.804		3.394		3.286	

Panel B: Estimation results with ARMA(2,1) model

	GM (#3180)		IBM (#3178)		MRK (#3186)	
Likelihood	-5191.41		-4505.64		-4493.27	
	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
σ^2	0.0163	(0.0060)	0.0341	(0.0150)	0.0807	(0.0321)
AR_2	-0.1300	(0.0485)	-0.2053	(0.0437)	-0.2350	(0.0662)
AR_1	0.1979	(0.0458)	0.3081	(0.0419)	0.3566	(0.0553)
MA_1	0.9269	(0.0130)	0.8821	(0.0281)	0.8428	(0.0460)
a	2.8013	(0.1505)	2.7383	(0.3184)	2.1510	(0.1583)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	22.126	[0.334]	15.441	[0.751]	13.052	[0.875]
$Q^2(20)$	2.102	[1.000]	1.474	[1.000]	0.669	[1.000]
$ARCH1-5$	0.206	[0.960]	0.242	[0.944]	0.066	[0.997]
Schwarz	3.266		2.837		2.822	

Notes: This table reports the estimation results for two models for RV. Panel A reports the results for the simple Constant Volatility case whereas Panel B reports the result of a more general ARMA(2,1) specification for the dyamics. Thus, the latter model allows for time varying volatility. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. In square brakets below all test statistics p-values are reported. The last row reports the Schwarz Information Criteria.

Table 3: Regression results for the unconditional model

Panel A: General Motors

Model	Gaussian		NIG(a,0)		NIG(a,b)	
Likelihood	−15482.10		−15005.97		−14998.57	
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0152	(0.0109)	−0.0050	(0.0096)	0.0152	(0.0110)
ω	2.3043	(0.0867)	2.2565	(0.0582)	2.2586	(0.0584)
a			1.0643	(0.0977)	1.0423	(0.0493)
b					0.0874	(0.0251)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	34648.895	(0.000)	34648.895	(0.000)	34648.895	(0.000)
$Q(20)$	65.505	(0.000)	65.505	(0.000)	65.505	(0.000)
$Q^2(20)$	1070.059	(0.000)	1070.053	(0.000)	1070.059	(0.000)
Schwarz	3.6728		3.5600		3.5582	*

Panel B: International Business Machines

Model	Gaussian		NIG(a,0)		NIG(a,b)	
Likelihood	−14963.52		−14441.62		−14436.15	
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0136	(0.0109)	−0.0021	(0.0099)	0.0136	(0.0111)
ω	2.0376	(0.0968)	1.9433	(0.0528)	1.9452	(0.0530)
a			1.2657	(0.1347)	1.1316	(0.0622)
b					0.0826	(0.0304)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	102110.830	(0.000)	102110.830	(0.000)	102110.830	(0.000)
$Q(20)$	42.111	(0.003)	42.111	(0.003)	42.111	(0.003)
$Q^2(20)$	390.498	(0.000)	390.572	(0.000)	390.498	(0.000)
Schwarz	3.5498		3.4261		3.4248	*

Panel C: Merk and Company Inc.

Model	Gaussian		NIG(a,0)		NIG(a,b)	
Likelihood	−15087.68		−14794.87		−14786.80	
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0372	(0.0109)	0.0186	(0.0099)	0.0373	(0.0109)
ω	2.0985	(0.0518)	2.0919	(0.0472)	2.0943	(0.0474)
a			1.3616	(0.1253)	1.1775	(0.0552)
b					0.1033	(0.0272)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	3480.030	(0.000)	3480.030	(0.000)	3480.030	(0.000)
$Q(20)$	108.552	(0.000)	108.552	(0.000)	108.552	(0.000)
$Q^2(20)$	1163.183	(0.000)	1154.877	(0.000)	1163.183	(0.000)
Schwarz	3.5793		3.5099		3.5080	*

Notes: This table reports the regression results for the unconditional models discussed in Section 3.

Table 4: Estimation results for General Motors (GM)

Panel A: Gaussian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	−9360.67		−9337.06		−9338.45	
	Estimate	Std.Err.	Estimat	Std.Err.	Estimate	Std.Err.
λ	0.0322	(0.0140)	0.0134	(0.0140)	0.0187	(0.0139)
ω	0.0305	(0.0207)	0.0191	(0.0102)	−0.0680	(0.0208)
β	0.9330	(0.0303)	0.9392	(0.0196)	0.9893	(0.0057)
α	0.0574	(0.0258)	0.0417	(0.0149)	0.1013	(0.0334)
γ/θ			−0.5668	(0.1233)	−0.4175	(0.0969)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	1815.684	[0.000]	1262.570	[0.000]	1229.000	[0.000]
$Q(20)$	34.200	[0.025]	32.802	[0.035]	30.902	[0.056]
$Q^2(20)$	37.661	[0.004]	36.893	[0.005]	43.758	[0.001]
$ARCH1-5$	4.344	[0.001]	3.426	[0.004]	4.913	[0.000]
$Schwarz$	3.704		3.695	*	3.695	

Panel B: Normal Inverse Gasssian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	−9242.60	(118.06)	−9232.51	(104.55)	−9232.00	(106.45)
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0214	(0.0138)	0.0129	(0.0141)	0.0169	(0.0140)
ω	0.0162	(0.0073)	0.0125	(0.0059)	−0.0507	(0.0124)
β	0.9584	(0.0103)	0.9571	(0.0094)	0.9938	(0.0031)
α	0.0359	(0.0088)	0.0307	(0.0076)	0.0734	(0.0190)
γ/θ			−0.5089	(0.1324)	−0.3817	(0.1075)
a	2.1707	(0.3396)	2.3072	(0.3620)	2.2722	(0.3508)
b	0.1756	(0.0779)	0.1898	(0.0835)	0.1967	(0.0831)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	33.712	[0.028]	32.628	[0.037]	30.544	[0.062]
$Q^2(20)$	65.639	[0.000]	55.436	[0.000]	81.257	[0.000]
$ARCH1-5$	10.620	[0.000]	7.938	[0.000]	12.966	[0.000]
$Schwarz$	3.657		3.654		3.653	*

Notes: This table reports Quasi Maximum Likelihood Estimates (QMLE) for the daily returns from 1976 through 1995 assuming a risk-free interest rate of 5.4% corresponding to the value on December 29, 1995. Robust standard errors are reported in parentheses. J-B is the value of the usual Jarque-Bera normality test for the standardized residuals. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. In square brakets below all test statistics p-values are reported. The last row reports the Schwarz Information Criteria, with an asterix denoting the minimum value.

Table 5: Estimation results for International Business Machines (IBM)

Panel A: Gaussian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	−8798.75		−8769.71		−8751.84	
	Estimate	Std.Err.	Estimat	Std.Err.	Estimate	Std.Err.
λ	0.0316	(0.0171)	0.0100	(0.0140)	0.0137	(0.0150)
ω	0.0246	(0.0143)	0.0270	(0.0131)	−0.0750	(0.0202)
β	0.9391	(0.0266)	0.9265	(0.0271)	0.9877	(0.0054)
α	0.0522	(0.0258)	0.0487	(0.0180)	0.1104	(0.0306)
γ/θ			−0.5433	(0.1230)	−0.3984	(0.0936)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	13673.471	[0.000]	7077.476	[0.000]	6122.562	[0.000]
$Q(20)$	24.823	[0.208]	24.900	[0.205]	26.107	[0.162]
$Q^2(20)$	9.161	[0.956]	10.164	[0.926]	12.175	[0.838]
$ARCH1-5$	0.441	[0.820]	0.392	[0.855]	0.894	[0.484]
Schwarz	3.482		3.470		3.463	*

Panel B: Normal Inverse Gasssian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	−8569.95	(228.80)	−8566.42	(203.29)	−8556.75	(195.09)
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0138	(0.0141)	0.0094	(0.0143)	0.0109	(0.0144)
ω	0.0221	(0.0078)	0.0227	(0.0081)	−0.0580	(0.0104)
β	0.9552	(0.0099)	0.9514	(0.0109)	0.9895	(0.0038)
α	0.0334	(0.0075)	0.0339	(0.0073)	0.0849	(0.0159)
γ/θ			−0.2954	(0.1144)	−0.2462	(0.0758)
a	1.7385	(0.2634)	1.7895	(0.2635)	1.8258	(0.2645)
b	0.1340	(0.0693)	0.1372	(0.0705)	0.1452	(0.0720)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	26.438	[0.152]	26.163	[0.161]	28.137	[0.106]
$Q^2(20)$	9.817	[0.938]	9.796	[0.938]	16.002	[0.592]
$ARCH1-5$	0.986	[0.424]	0.842	[0.520]	2.059	[0.067]
Schwarz	3.391		3.390		3.386	*

Notes: See the notes to Table 4.

Table 6: Estimation results for Merck & Company Inc. (MRK)

Panel A: Gaussian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	-8858.53		-8849.72		-8847.33	
	Estimate	Std.Err.	Estimat	Std.Err.	Estimate	Std.Err.
λ	0.0512	(0.0139)	0.0407	(0.0141)	0.0415	(0.0142)
ω	0.0692	(0.0281)	0.0628	(0.0198)	-0.0748	(0.0152)
β	0.9072	(0.0276)	0.9100	(0.0205)	0.9690	(0.0097)
α	0.0600	(0.0169)	0.0535	(0.0128)	0.1253	(0.0256)
γ/θ			-0.3505	(0.1211)	-0.2621	(0.0894)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$J-B$	523.776	[0.000]	431.803	[0.000]	429.685	[0.000]
$Q(20)$	38.847	[0.007]	39.009	[0.007]	37.890	[0.009]
$Q^2(20)$	20.024	[0.331]	22.846	[0.197]	23.243	[0.181]
$ARCH1-5$	1.434	[0.208]	1.807	[0.108]	1.939	[0.085]
$Schwarz$	3.505		3.502		3.501	*

Panel B: Normal Inverse Gasssian models

Model	GARCH		NGARCH		EGARCH	
Likelihood	-8762.44	(96.09)	-8757.27	(92.45)	-8755.19	(92.15)
	Estimate	Std.Err.	Estimate	Std.Err.	Estimate	Std.Err.
λ	0.0469	(0.0138)	0.0406	(0.0140)	0.0410	(0.0140)
ω	0.0607	(0.0241)	0.0585	(0.0181)	-0.0701	(0.0124)
β	0.9188	(0.0221)	0.9170	(0.0169)	0.9722	(0.0087)
α	0.0525	(0.0124)	0.0496	(0.0099)	0.1168	(0.0208)
γ/θ			-0.3346	(0.1133)	-0.2459	(0.0813)
a	2.0417	(0.2736)	2.0827	(0.2774)	2.0884	(0.2787)
b	0.2191	(0.0703)	0.2205	(0.0716)	0.2196	(0.0715)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	38.571	[0.008]	38.813	[0.007]	37.749	[0.010]
$Q^2(20)$	21.684	[0.246]	23.304	[0.179]	24.447	[0.141]
$ARCH1-5$	1.858	[0.098]	2.010	[0.074]	2.183	[0.053]
$Schwarz$	3.467		3.466		3.465	*

Notes: See the notes to Table 4.

Table 7: American price estimates in a model with GARCH volatility processes

Panel A: American put price estimates

		Gaussian		Symmetric NIG			NIG		
T	Strike	Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	85	0.000	(0.0002)	0.001	(0.0005)	−73.141%	0.001	(0.0004)	−55.93%
7	100	1.556	(0.0100)	1.535	(0.0110)	1.425%	1.529	(0.0110)	1.78%
7	115	15.000	(0.0016)	15.000	(0.0013)	0.000%	15.000	(0.0010)	0.00%
21	85	0.036	(0.0027)	0.047	(0.0041)	−23.466%	0.037	(0.0034)	−3.81%
21	100	2.648	(0.0183)	2.612	(0.0213)	1.383%	2.600	(0.0212)	1.84%
21	115	15.005	(0.0093)	15.005	(0.0080)	0.001%	15.007	(0.0098)	−0.01%
63	85	0.473	(0.0151)	0.491	(0.0190)	−3.572%	0.441	(0.0168)	7.19%
63	100	4.454	(0.0382)	4.401	(0.0388)	1.218%	4.371	(0.0381)	1.90%
63	115	15.323	(0.0389)	15.301	(0.0372)	0.147%	15.328	(0.0375)	−0.03%
126	85	1.327	(0.0361)	1.329	(0.0456)	−0.172%	1.241	(0.0422)	6.87%
126	100	6.141	(0.0545)	6.060	(0.0619)	1.329%	6.008	(0.0633)	2.21%
126	115	16.122	(0.0612)	16.057	(0.0609)	0.404%	16.080	(0.0630)	0.26%

Panel B: American call price estimates

		Gaussian		Symmetric NIG			NIG		
T	Strike	Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	115	0.001	(0.0004)	0.003	(0.0010)	−58.873%	0.005	(0.0012)	−71.10%
7	100	1.637	(0.0115)	1.616	(0.0130)	1.313%	1.611	(0.0130)	1.62%
7	85	15.071	(0.0080)	15.071	(0.0078)	0.005%	15.071	(0.0080)	0.00%
21	115	0.090	(0.0053)	0.101	(0.0070)	−11.143%	0.115	(0.0079)	−22.14%
21	100	2.883	(0.0252)	2.851	(0.0261)	1.122%	2.840	(0.0261)	1.51%
21	85	15.223	(0.0135)	15.229	(0.0108)	−0.038%	15.223	(0.0117)	0.00%
63	115	0.933	(0.0257)	0.926	(0.0285)	0.760%	0.964	(0.0309)	−3.26%
63	100	5.150	(0.0495)	5.101	(0.0487)	0.954%	5.073	(0.0494)	1.51%
63	85	15.986	(0.0251)	16.003	(0.0281)	−0.105%	15.955	(0.0283)	0.19%
126	115	2.501	(0.0605)	2.465	(0.0688)	1.489%	2.499	(0.0700)	0.10%
126	100	7.488	(0.0765)	7.426	(0.0870)	0.844%	7.376	(0.0863)	1.53%
126	85	17.327	(0.0533)	17.332	(0.0613)	−0.031%	17.247	(0.0602)	0.46%

Notes: This table shows American option prices for a GARCH variance specification with different underlying distributions for a set of artificial options. The interest rate is fixed at 6% with a dividend yield of 3% both of which are annualized using 252 days a year. T denotes the time to maturity in days, Strike denotes the strike price, and for all the options the stock price is 100. The parameter values for the different GARCH processes are the ones specified in the text. In the cross-sectional regressions powers of and cross-products between the stock level and the volatility level of total order less than or equal to two were used. Exercise is considered once every trading day. Prices reported are averages of 100 calculated prices using 20.000 paths and different seeds in the random number generator, "rann", in Ox. In parentheses standard errors of these price estimates are reported. The columns headed Rel Diff reports the difference between the Gaussian GARCH specification and corresponding NIG GARCH specification. Thus, they indicate the mispricing which would occur by the Gaussian GARCH volatility model if the true model is the corresponding NIG GARCH model.

Table 8: American price estimates in a model with NGARCH volatility processes

Panel A: American put price estimates

		Gaussian		Symmetric NIG			NIG		
T	Strike	Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	85	0.001	(0.0003)	0.002	(0.0007)	−62.766%	0.001	(0.0005)	−42.03%
7	100	1.613	(0.0110)	1.591	(0.0121)	1.358%	1.587	(0.0121)	1.59%
7	115	15.000	(0.0014)	15.000	(0.0002)	0.002%	15.000	(0.0010)	0.00%
21	85	0.068	(0.0041)	0.078	(0.0058)	−12.849%	0.066	(0.0050)	2.88%
21	100	2.751	(0.0203)	2.718	(0.0231)	1.218%	2.710	(0.0231)	1.50%
21	115	15.003	(0.0060)	15.001	(0.0040)	0.009%	15.003	(0.0054)	0.00%
63	85	0.689	(0.0194)	0.690	(0.0239)	−0.257%	0.645	(0.0212)	6.84%
63	100	4.649	(0.0403)	4.599	(0.0409)	1.070%	4.583	(0.0395)	1.42%
63	115	15.259	(0.0377)	15.248	(0.0342)	0.070%	15.271	(0.0369)	−0.08%
126	85	1.724	(0.0418)	1.704	(0.0506)	1.153%	1.631	(0.0480)	5.67%
126	100	6.446	(0.0556)	6.365	(0.0692)	1.271%	6.340	(0.0643)	1.67%
126	115	16.097	(0.0572)	16.037	(0.0645)	0.372%	16.069	(0.0631)	0.17%

Panel B: American call price estimates

		Gaussian		Symmetric NIG			NIG		
T	Strike	Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	115	0.001	(0.0003)	0.002	(0.0007)	−67.688%	0.003	(0.0009)	−78.37%
7	100	1.692	(0.0111)	1.672	(0.0128)	1.217%	1.668	(0.0128)	1.41%
7	85	15.073	(0.0086)	15.073	(0.0081)	0.001%	15.071	(0.0081)	0.01%
21	115	0.064	(0.0040)	0.075	(0.0053)	−15.070%	0.087	(0.0061)	−27.36%
21	100	2.985	(0.0232)	2.956	(0.0251)	0.986%	2.949	(0.0251)	1.21%
21	85	15.246	(0.0117)	15.257	(0.0105)	−0.067%	15.244	(0.0110)	0.01%
63	115	0.843	(0.0219)	0.845	(0.0255)	−0.325%	0.886	(0.0269)	−4.93%
63	100	5.344	(0.0458)	5.300	(0.0463)	0.828%	5.286	(0.0463)	1.09%
63	85	16.196	(0.0202)	16.199	(0.0256)	−0.016%	16.155	(0.0250)	0.26%
126	115	2.458	(0.0541)	2.433	(0.0622)	1.034%	2.479	(0.0629)	−0.85%
126	100	7.786	(0.0714)	7.727	(0.0809)	0.762%	7.701	(0.0795)	1.10%
126	85	17.715	(0.0485)	17.700	(0.0568)	0.085%	17.629	(0.0552)	0.49%

Notes: This table shows American put prices and early exercise premiums for a NGARCH variance specification with different underlying distributions for the set of artificial options in Table 7 using the parameter values from the text.

Table 9: Parameter Estimates

Parameter estimates from "implied" estimation:

	GM		IBM		MRK	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0234	(0.0130)	0.0163	(0.0128)	0.0449	(0.0131)
ω	0.0152	(0.0070)	0.0210	(0.0076)	0.0610	(0.0238)
β	0.9596	(0.0100)	0.9564	(0.0097)	0.9191	(0.0217)
α	0.0349	(0.0085)	0.0326	(0.0073)	0.0517	(0.0121)
a	2.2124	(0.3544)	1.7487	(0.2692)	2.0830	(0.2814)

Parameter estimates from "regular" estimation:

	GM		IBM		MRK	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0235	(0.0133)	0.0166	(0.0132)	0.0455	(0.0134)
ω	0.0152	(0.0070)	0.0210	(0.0076)	0.0611	(0.0238)
β	0.9596	(0.0100)	0.9564	(0.0097)	0.9190	(0.0217)
α	0.0349	(0.0085)	0.0326	(0.0073)	0.0517	(0.0121)
a	2.2121	(0.3500)	1.7490	(0.2623)	2.0824	(0.2832)

Notes: This table reports parameter estimates for GARCH(1,1) models with Normal Inverse Gaussian errors.

Table 10: Overall performance for the NIG GARCH option pricing model

Panel A: All stocks

Model	All	(8424)	Put	(3009)	Call	(5415)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	−20.34%	10.86%	−22.90%	12.21%	−18.91%	10.10%
GARCH	−11.12%	6.84%	−13.75%	7.77%	−9.66%	6.33%
NGARCH	−11.20%	6.80%	−11.87%	7.10%	−10.83%	6.64%
EGARCH	−12.34%	6.00%	−12.75%	6.17%	−12.12%	5.90%
Best	2	4	3	4	2	4
Worst	1	1	1	1	1	1

Panel B: GM

Model	All	(1835)	Put	(629)	Call	(1206)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	−17.45%	7.46%	−20.47%	9.81%	−15.88%	6.24%
GARCH	−2.47%	3.28%	−4.29%	4.36%	−1.53%	2.72%
NGARCH	−2.49%	3.07%	−1.81%	3.92%	−2.84%	2.62%
EGARCH	−2.83%	2.62%	−2.28%	3.29%	−3.12%	2.27%
Best	2	4	3	4	2	4
Worst	1	1	1	1	1	1

Panel C: IBM

Model	All	(4745)	Put	(1827)	Call	(2918)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	−24.52%	13.29%	−25.79%	14.03%	−23.72%	12.83%
GARCH	−14.70%	8.82%	−16.83%	9.41%	−13.36%	8.45%
NGARCH	−14.94%	8.97%	−15.21%	8.79%	−14.76%	9.08%
EGARCH	−16.94%	7.66%	−16.60%	7.47%	−17.14%	7.78%
Best	2	4	3	4	2	4
Worst	1	1	1	1	1	1

Panel D: MRK

Model	All	(1844)	Put	(553)	Call	(1291)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	−12.45%	7.98%	−16.11%	8.96%	−10.88%	7.55%
GARCH	−10.53%	5.29%	−14.33%	6.23%	−8.90%	4.89%
NGARCH	−10.25%	4.95%	−12.26%	5.15%	−9.39%	4.86%
EGARCH	−9.99%	5.07%	−11.90%	5.15%	−9.17%	5.04%
Best	4	3	4	4	2	3
Worst	1	1	1	1	1	1

Notes: This table reports results on the overall pricing performance of the NIG GARCH option pricing model using the metrics in Section 6.

Table 11: Pricing performance for the NIG GARCH option pricing model relative to the Gaussian GARCH modell

Panel A: All stocks

Model	All	(8424)	Put	(3009)	Call	(5415)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.157	1.103	1.179	1.138	1.142	1.081
GARCH	1.771	0.867	1.781	0.998	1.763	0.796
NGARCH	1.365	0.862	2.269	0.962	1.098	0.812
EGARCH	2.083	1.050	3.404	1.159	1.698	0.995
Max	2.083	1.103	3.404	1.159	1.763	1.081
Min	1.157	0.862	1.179	0.962	1.098	0.796

Panel B: GM

Model	All	(1835)	Put	(629)	Call	(1206)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.106	1.082	1.126	1.114	1.093	1.056
GARCH	1.490	0.894	1.710	0.975	1.253	0.837
NGARCH	1.055	0.870	2.994	0.944	0.728	0.820
EGARCH	18.408	0.897	1.175	0.927	2.496	0.875
Max	18.408	1.082	2.994	1.114	2.496	1.056
Min	1.055	0.870	1.126	0.927	0.728	0.820

Panel C: IBM

Model	All	(4745)	Put	(1827)	Call	(2918)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.198	1.135	1.219	1.169	1.185	1.112
GARCH	2.202	0.844	2.099	0.992	2.292	0.765
NGARCH	1.517	0.844	2.678	0.947	1.185	0.791
EGARCH	2.131	1.075	3.629	1.206	1.705	1.009
Max	2.202	1.135	3.629	1.206	2.292	1.112
Min	1.198	0.844	1.219	0.947	1.185	0.765

Panel D: MRK

Model	All	(1844)	Put	(553)	Call	(1291)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.040	1.002	1.066	1.024	1.025	0.991
GARCH	1.067	0.964	1.132	1.050	1.026	0.923
NGARCH	1.047	0.953	1.180	1.079	0.985	0.906
EGARCH	1.544	1.047	1.593	1.154	1.519	1.006
Max	1.544	1.047	1.593	1.154	1.519	1.006
Min	1.040	0.953	1.066	1.024	0.985	0.906

Notes: This table reports the pricing results of the NIG GARCH model relative to what is obtained when the Gaussian GARCH model is used for each of the respective models using the metrics in Section 6. Thus, in the table a cell value smaller than one indicates that for this model the NIG GARCH specification has smaller pricing errors than the Gaussian GARCH model. If a cell value larger than one indicates that the Gaussian model had smallest pricing errors.

Table 12: Pricing performance for the NIG GARCH option pricing model relative to the Gaussian GARCH model in terms of moneyness and maturity

Panel A: Oveall performance in terms of moneyness

Model	ITM	(2609)	ATM	(1751)	OTM	(4064)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	-6.54%	1.03%	-9.84%	3.37%	-33.72%	20.39%
GARCH	-3.45%	0.57%	-6.01%	2.33%	-18.25%	12.81%
NGARCH	-3.12%	0.53%	-5.94%	2.27%	-18.65%	12.79%
EGARCH	-3.40%	0.46%	-6.05%	1.92%	-20.79%	11.31%
Best	3	4	3	4	2	4
Worst	1	1	1	1	1	1

Panel B: Relative performance in terms of moneyness

Model	ITM	(2609)	ATM	(1751)	OTM	(4064)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.161	1.164	1.364	1.131	1.134	1.099
GARCH	1.838	1.012	2.256	0.948	1.711	0.858
NGARCH	1.797	1.039	1.660	0.946	1.300	0.852
EGARCH	3.284	1.260	5.018	1.039	1.874	1.046
Max	3.284	1.260	5.018	1.131	1.874	1.099
Min	1.161	1.012	1.364	0.946	1.134	0.852

Panel C: Overall performance in terms of maturity

Model	ST	(4687)	MT	(3208)	LT	(529)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	-20.22%	11.95%	-20.69%	9.67%	-19.23%	8.36%
GARCH	-10.76%	7.92%	-11.31%	5.57%	-13.23%	4.97%
NGARCH	-10.89%	7.93%	-11.28%	5.47%	-13.42%	4.94%
EGARCH	-11.78%	6.87%	-12.82%	4.92%	-14.42%	4.81%
Best	2	4	3	4	2	4
Worst	1	1	1	1	1	1

Panel D: Relative performance in terms of maturity

Model	ST	(4687)	MT	(3208)	LT	(529)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
CV	1.143	1.077	1.174	1.146	1.172	1.157
GARCH	1.503	0.827	2.257	0.923	2.131	1.201
NGARCH	1.277	0.827	1.486	0.913	1.480	1.109
EGARCH	1.493	0.970	3.538	1.189	5.840	1.527
Max	1.503	1.077	3.538	1.189	5.840	1.527
Min	1.143	0.827	1.174	0.913	1.172	1.109

Notes: This table reports the pricing results of the NIG GARCH model relative to what is obtained when the Gaussian GARCH model is used for each of the respective models using the metrics in Section 6. Thus, the reported results have the same interpretation as those in Table 11. The reported results are for all the stock options considered together but for different moneyness and maturity categories. Thus, ITM is in the money options, ATM is at the money options, and OTM is out of the money options. Likewise, ST refers to short term options, MT to middle term options, and LT to long term options.

Figures

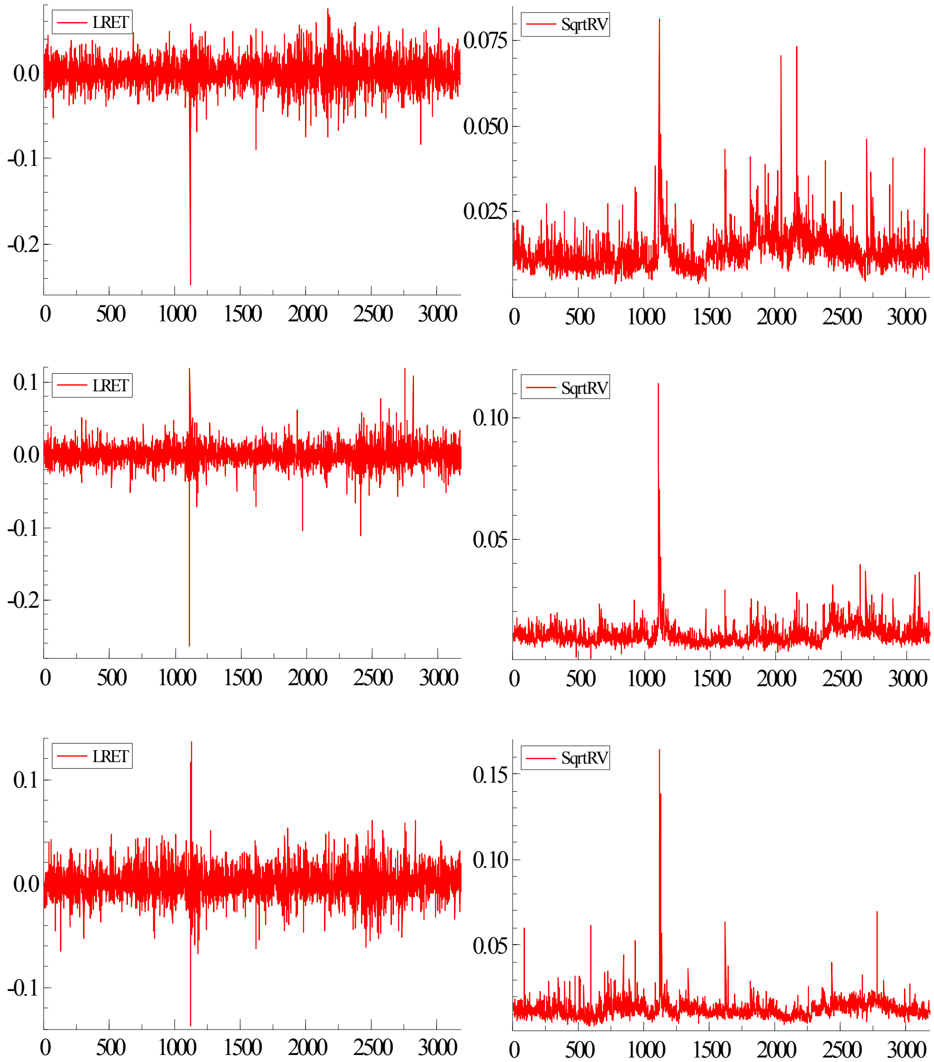

Figure 1: Time plots of daily returns, R_t, and realized volatilities, $\sqrt{RV_t}$.

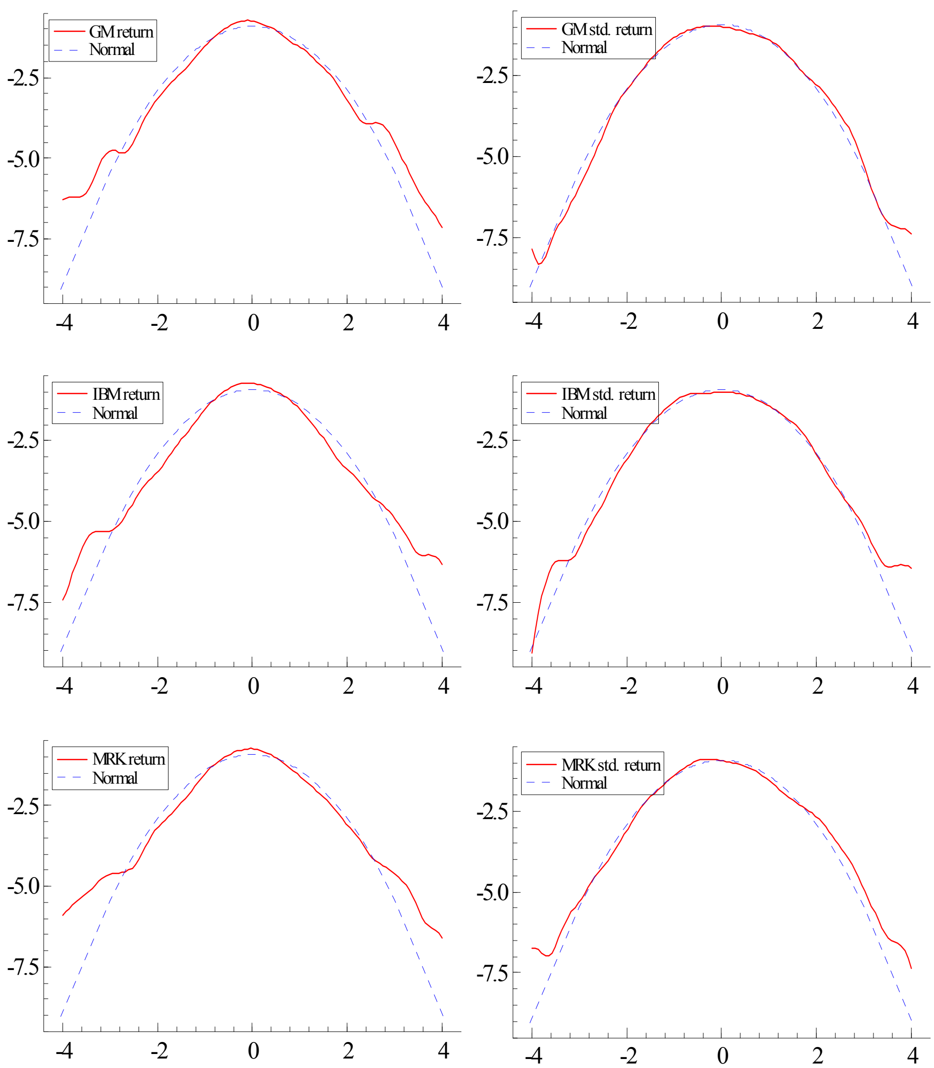

Figure 2: Log density plots of raw return series, R_t, and standardized return, $\frac{R_t}{\sqrt{RV_t}}$. The dotted line plots the normal density with the same mean and variance.

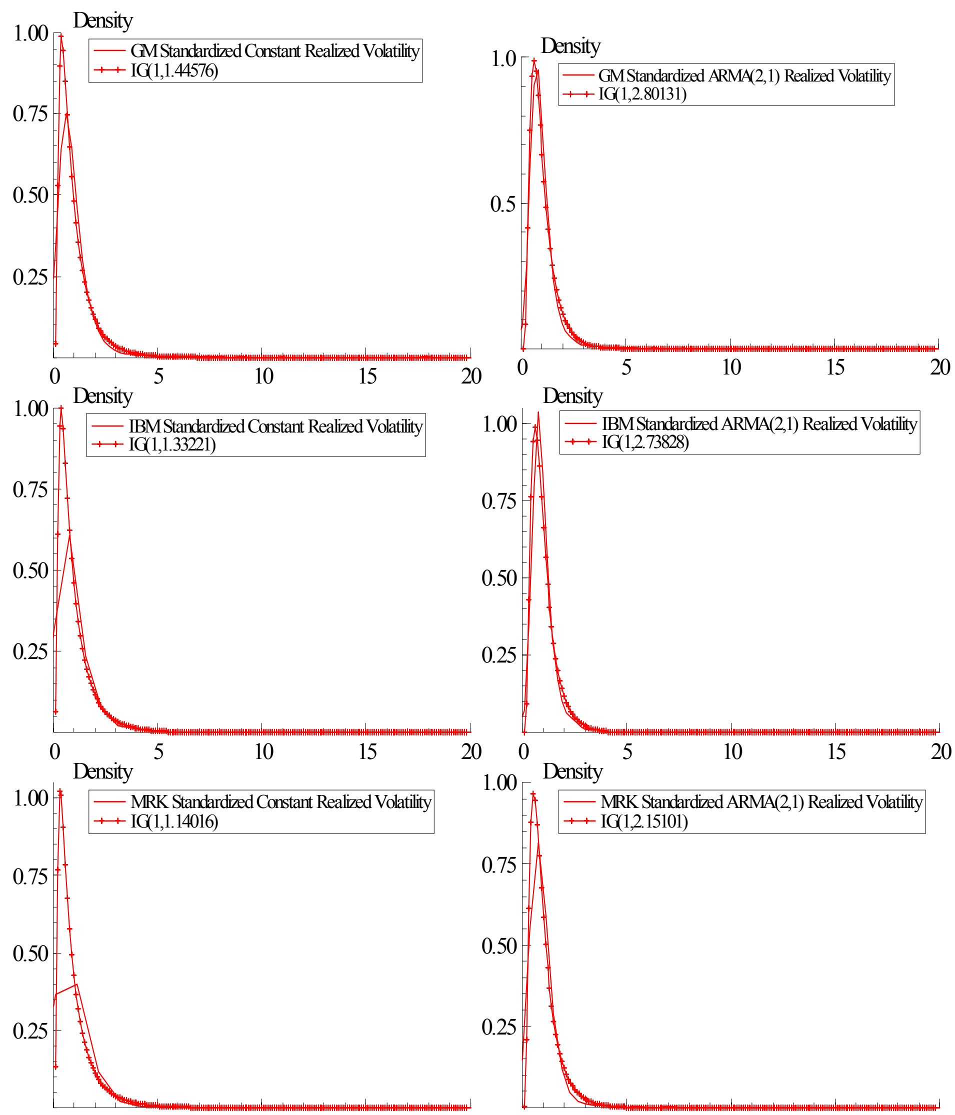

Figure 3: Plots of the residual desities for the two specifications of RV considered. Left hand panels show residuals from a simple specification with only a constant term, which corresponds to a constant volatility model. Right hand panels show residuals from a ARMA(2,1) specification whichs corresponds to a model with time varying volatility.

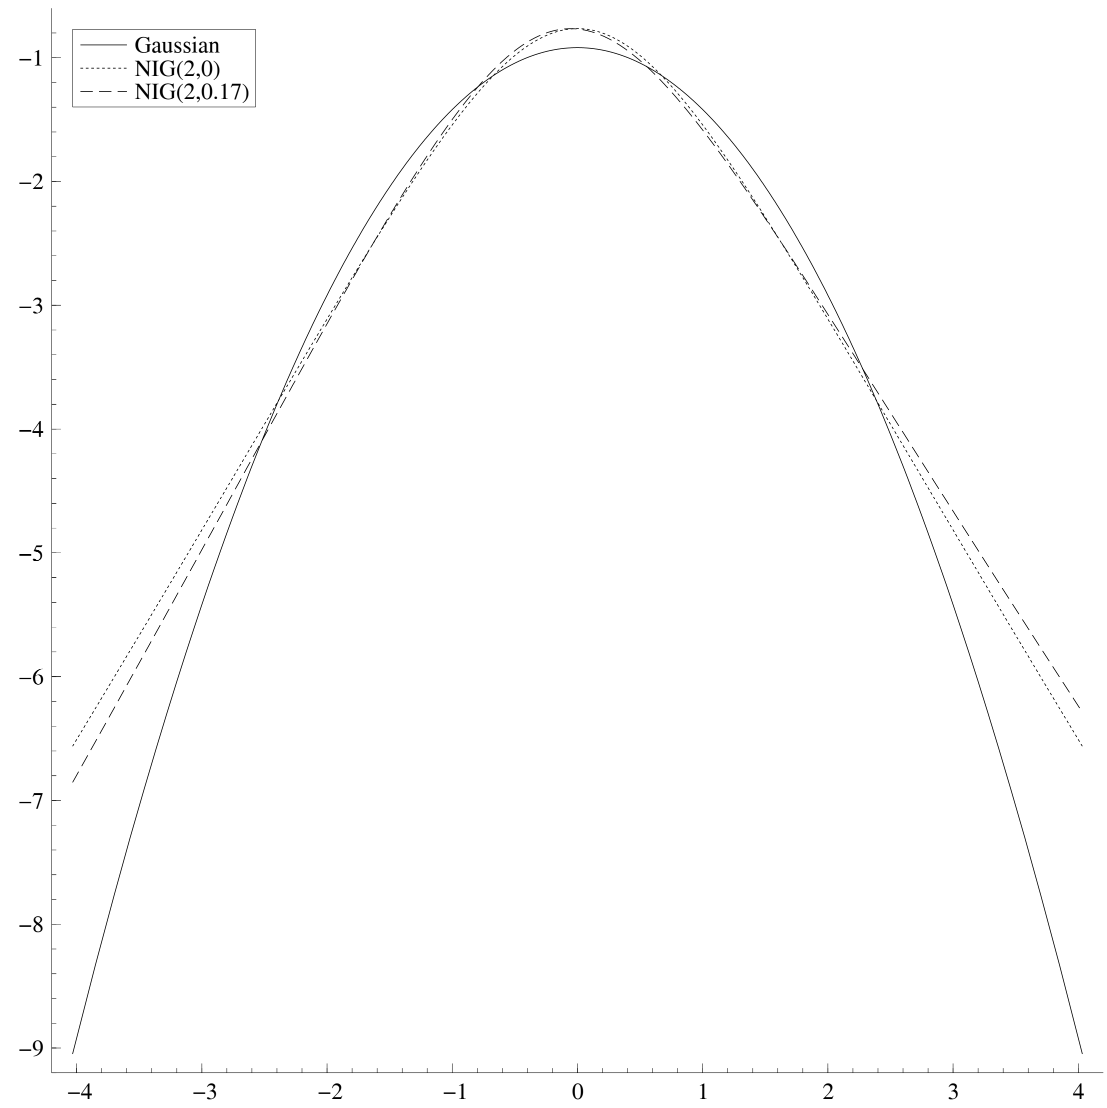

Figure 4: Plots of log densities for the Gaussian distribution, the symmetric NIG distribution with $a = 2$, and the NIG distribution with $a = 2$ and $b = 0.17$.

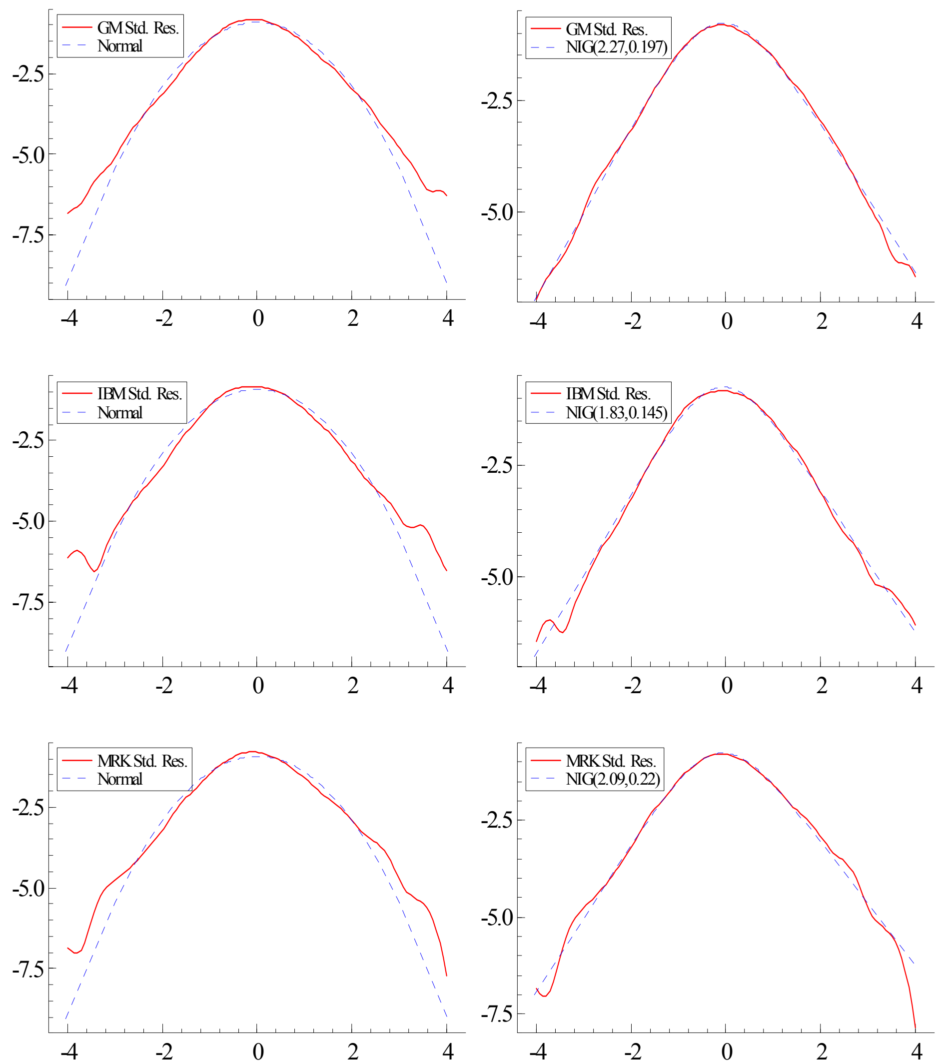

Figure 5: Plot of residual log-densities from a EGARCH volatility specification for General Motor (GM), International Business Machines (IBM), and Merck & Company Inc. (MRK). Left hand column shows plots for the Gaussian model and right hand column for the general NIG model.

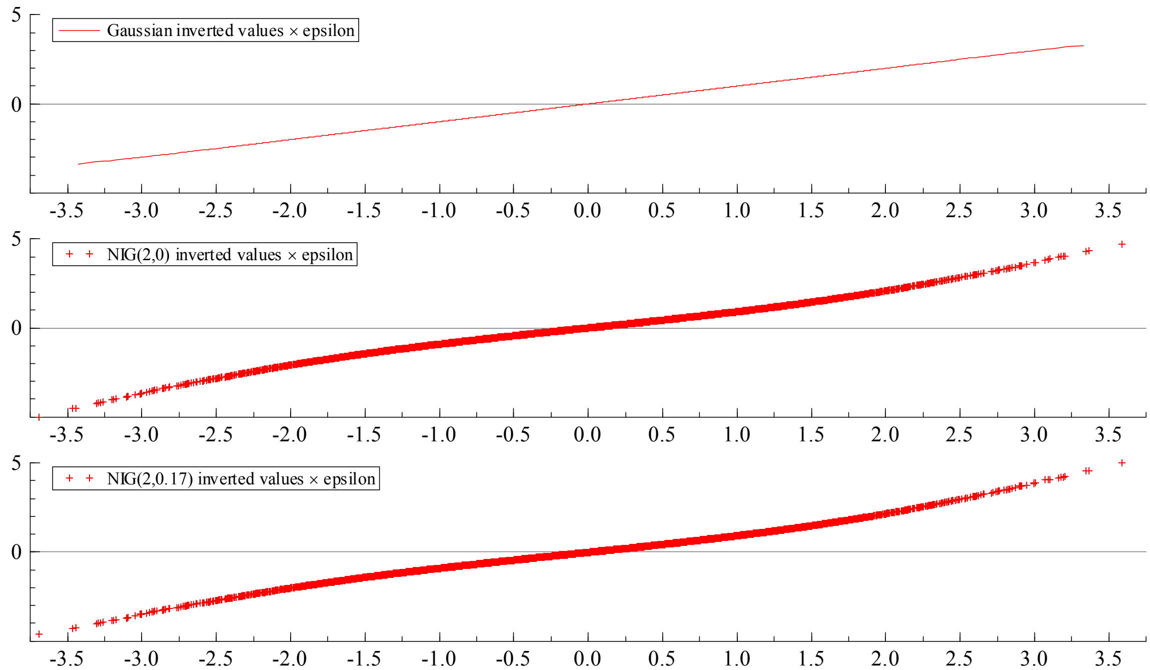

Figure 6: Cross-plots of the inverted values, $F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ with $\lambda = 0.05$ against $10,000$ normally distributed Z_t, for the Gaussian distribution, the symmetric NIG distribution with $a = 2$ and $b = 0$, and the NIG distribution with $a = 2$ and $b = 0.17$.

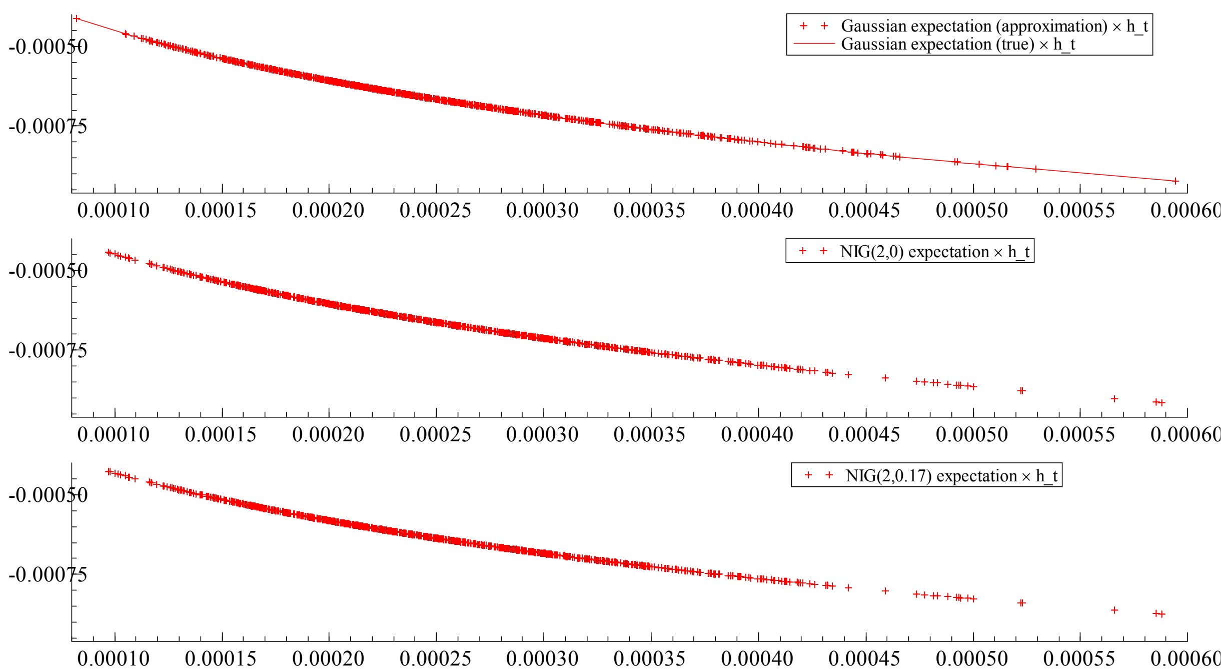

Figure 7: Cross plots of the logarithm of the expected values $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_D^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ against values of h_t for the Gaussian distribution, the symmetric NIG distribution with $a=2$ and $b=0$, and the NIG distribution with $a=2$ and $b=0.17$.

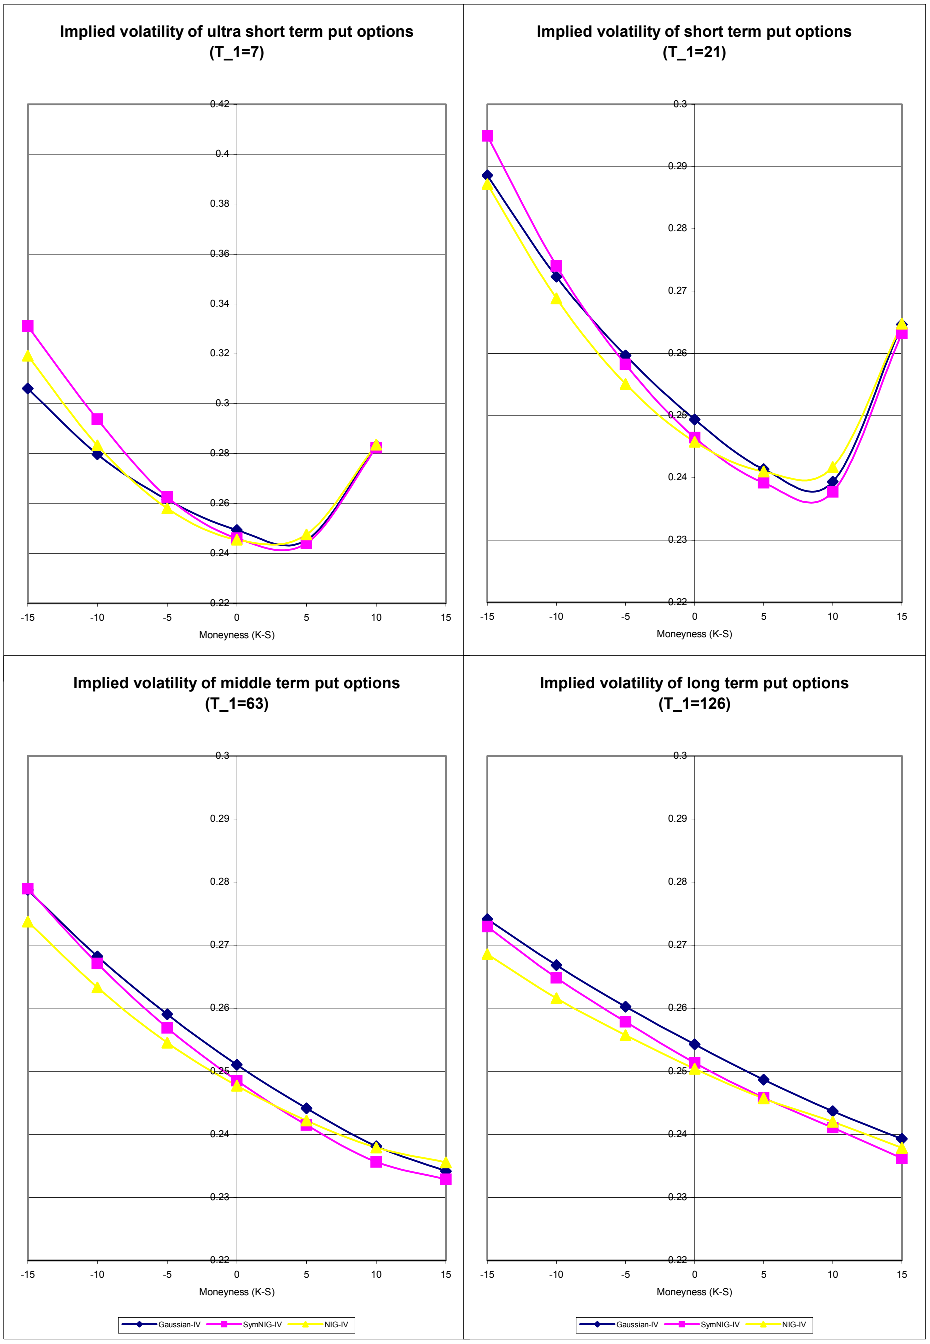

Figure 8: Implied volatilities for put options with the NGARCH volatility specification plotted against the difference between the strike price and the level with different underlying distributions.

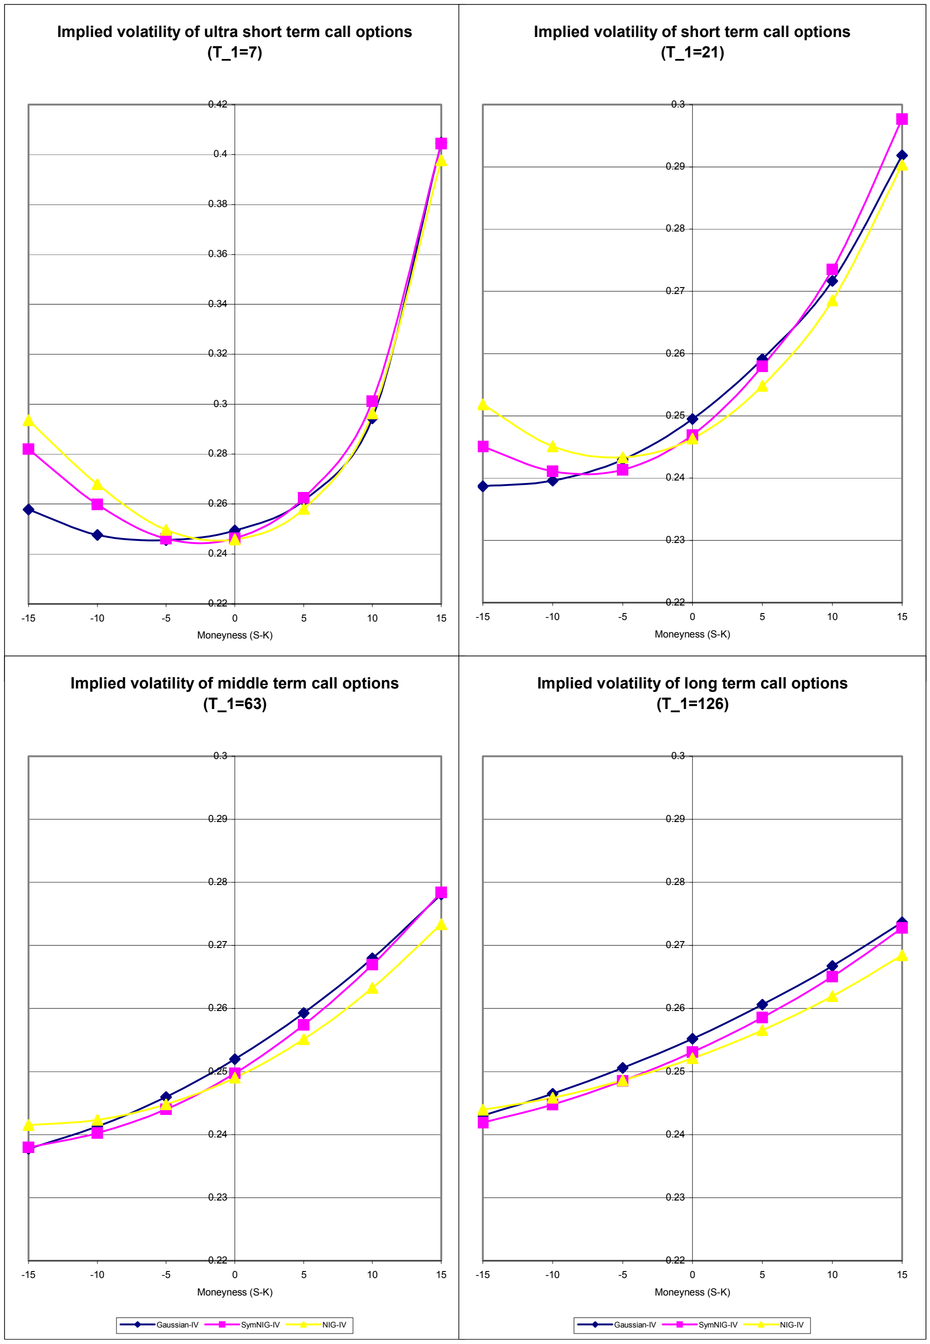

Figure 9: Implied volatilities for call options with the NGARCH volatility specification plotted against the difference between the strike price and the level with different underlying distributions.